Predictive Inference for Spatio-temporal Precipitation Data and Its Extremes*

Yang Liu† and Philip Kokic
CSIRO, Canberra, Australia

Version of October 17, 2014

arXiv:1411.4715v1 [stat.AP] 18 Nov 2014

Under review at the *Journal of the American Statistical Association*

Abstract

Modelling of precipitation and its extremes is important for urban and agriculture planning purposes. We present a method for producing spatial predictions and measures of uncertainty for spatio-temporal data that is heavy-tailed and subject to substaintial skewness which often arise in measurements of many environmental processes, and we apply the method to precipitation data in south-west Western Australia. A generalised hyperbolic Bayesian hierarchical model is constructed for the intensity, frequency and duration of daily precipitation, including the extremes. Unlike models based on extreme value theory, which only model maxima of finite-sized blocks or exceedances above a large threshold, the proposed model uses all the data available efficiently, and hence not only fits the extremes but also models the entire rainfall distribution. It captures spatial and temporal clustering, as well as spatially and temporally varying volatility and skewness. The model assumes that the regional precipitation is driven by a latent process characterised by geographical and climatological covariates. Effects not fully described by the covariates are captured by spatial and temporal structure in the hierarchies. Inference is provided by MCMC using a Metropolis-Hastings algorithm and spatial interpolation method, which provide a natural approach for estimating uncertainty. Similarly, both spatial and temporal predictions with uncertainty can be produced with the model.

Keywords: Bayesian inference, hierarchical modelling, non-Gaussian processes, spatial prediction.

1 Introduction

Statistical modelling of precipitation has important applications in many different fields of research including hydrology, agriculture, and environmental sciences. Within each of these fields information is required on a number of spatial and temporal scales. In hydrological applications the frequency and duration of extreme precipitation events over short time periods is very important (Shao et al., 2013; Fowler et al., 2007; Tetzlaff et al., 2005). For agricultural modelling climate information is required on the short-term variations associated with extreme and non-extreme events through the realistic simulation of observed data (Keating et al., 2003; Kokic et al., 2013), as well as intermediate variations associate with seasonal and inter-seasonal variations, and long-term variations due to causes such as climate change. Environmental science applications often require precipitation data with a high degree of spatial resolution (Ashcroft et al., 2011; Ashcroft and Gollan, 2012). For these reasons, statistical approaches that can provide consistent information covering these varying temporal and spatial characteristics is valuable. To address these requirements we need a flexible statistical approach that can accurately represent extreme events, and varying shapes and skewness of the rainfall distribution, as well as reliable estimates of the serial and spatial dependencies in

*The authors would like to thank Prof. Noel Cressie who provided insightful comments on this paper.
†Corresponding Author. Email: Yang.Liu@csiro.au

these data.

The objective of this paper is to describe a unified statistical model that moves towards meeting several of these objectives through the use of a Bayesian hierarchical model that utilises generalised hyperbolic (GH) processes. This method has significant advantages over models built upon statistical extreme value theory (EVT). The proposed model allows one to study both high and low precipitation events in a unified model. It is more data efficient, because it does not involve modelling maxima of finite-sized blocks or choosing a large threshold, so a much larger sample that may contain additional information will contribute to the estimates.

This research focuses on hydrometeorological data, but we note that our spatio-temporal model is not limited to this context, and this methodology can be adapted to other types of data and applications.

1.1 Measures for precipitation and extremes

Estimates of potential flooding and rainfall deficiency are necessary for city, rural area development planning and risk assessment. A commonly used measure of extreme events is the return period. It is a statistical measurement typically based on historic data denoting the average recurrence interval over an extended period of time. The calculation of return period assumes that the probability of the event occurring does not vary over time and is independent of past events. Take $p = 1/T = 1 - P(X \leq x)$, where T is the return interval, $P(X \leq x)$ is the cumulative distribution function of rainfall and x is an extreme event. Practically, for x large enough, the independence of events can be assumed. If the probability of an event occurring is p, then the probability of the event not occurring is $q = 1 - p$. The binomial distribution can be used to find the probability of occurrence of an event r times in a period of n years:

$$\binom{n}{r} \times p^r \times q^{n-r}. \tag{1}$$

Other frequently used uncertainty measures for precipitation and the extremes are duration over/below thresholds, as well as duration of zero precipitation. For aggregated data (eg. monthly and yearly), the number of time units over thresholds are more appropriate measures. These quantities provide probabilistic measures for prolonged extreme events and are highly appreciated in hydrological and environmental research.

1.2 Extreme value statistics

Rainfall arises from physical processes, but it is widely known that physical models such as General Circulation Models (GCMs) are inadequate for extremes, due to their coarse spatial resolution and the current incomplete understanding of the climate system (Ye and Li, 2011). Statistical models are therefore often considered when modelling extreme precipitation events.

Extreme value theory (EVT) is a frequently used approach for modelling extremes, because it provides statistical models for the tail of a probability distribution and complements to modeling the mean or central part of a distribution. EVT is based on the asymptotic arguments that lead to the generalized extreme value (GEV) distribution. For simplicity, in the univariate case, given independent and indentically distributed continuous data $Z_1, Z_2, \ldots, Z_n$, and letting $M_n = \max(Z_1, Z_2, \ldots, Z_n)$, it is known that if the normalized distribution of M_n converges as $n \to \infty$, then it converges to a GEV distribution (Fisher and Tippett, 1928; Gnedenko, 1943; Von Mises, 1936). Because of its asymptotic justification, the GEV distribution is used to model maxima of finite-sized blocks such as annual maxima. For example, this method was used by Nadarajah and Choi (2007) and Feng et al. (2007) for modelling annual rainfall maxima in China and

South Korea respectively. The model was fitted to individual weather stations, and the spatial distribution of the extreme rainfall return period was calculated through spatial interpolation. A more sophisticated model was used by Gaetan and Grigoletto (2007). The authors analysed annual rainfall maxima at weather stations in northeastern Italy using nonstationary spatial dependence and a random temporal trend in the parameters of the GEV distribution.

However, when daily precipitation is available, models which only use each year's annual maximum disregard other extreme events. The Generalized Pareto (GP) distribution is based on the exceedances above a threshold (Pickands III, 1975). Exceedances (the amounts which observations exceed a threshold u) should approximately follow a GP distribution as u gets large and the sample size increases. In this case, the tail of the distribution is characterized by the equation

$$P\left(Z > z + u | Z > u\right) = \max\left(0, \left(1 + \kappa \frac{z}{\sigma}\right)^{-1/\kappa}\right). \tag{2}$$

The scale parameter σ must be greater than zero, and the shape parameter κ controls whether the tail is bounded ($\kappa < 0$), light ($\kappa \to 0$), or heavy ($\kappa > 0$). The GP approach has been widely applied in the literature. Cooley et al. (2007) used the method with a common threshold at all weather stations to map return levels for extreme precipitation in Colorado. A stationary isotropic exponential covariance function was used to induce spatial dependence for these parameters. The shape parameter κ had two values depending on the station's location. Turkman et al. (2010) contructed a similar but more complex model for space-time properties of wildfires in Portugal, using a random walk to describe the temporal properties, and smoothing for the spatial modelling of exceedances. Love (2012) and Yuen and Guttorp (2014) also used GP spatial models for similar environmental data.

In practice, a threshold is chosen at a level where the data above it approximately follows a GP distribution and the shape and scale parameters are estimated. In order for the exceedances to follow a GP distribution, the chosen threshold is often large, so an enormous amount of data that could potentially provide additional information is discarded. In addition, threshold selection can be rather subjective. It is usually done using diagnostic plots that show how quantities such as the shape parameter vary as the threshold changes. Once chosen, the uncertainty associated with the choice of threshold is not accounted for (Coles and Tawn, 1996). Furthermore, Cooley et al. (2007) argued that the low-precision of rainfall data can also introduce a problem with threshold selection.

Ideally we would like to develop a more flexible spatio-temporal model which overcomes these problems. To this end, we need a class of distributions which can approximate the power law decay of the GP distribution in the tails, as well as provide flexibility in the centre and shoulder of the distribution. To meet these requirements we consider the generalised hyperbolic (GH) distribution.

1.3 The generalised hyperbolic distribution

The GH distribution was first introduced by Barndorff-Nielsen (1977) in connection with dune movements modelling. The Lebesgue density of the GH distribution is defined as

$$d_{GH(\lambda,\chi,\psi,\boldsymbol{\mu},\boldsymbol{\Sigma},\boldsymbol{\gamma})}(\boldsymbol{x}) = \frac{\left(\sqrt{\frac{\psi}{\chi}}\right)^{\lambda}\left(\psi+\boldsymbol{\gamma}^2\boldsymbol{\Sigma}^{-1}\right)^{\frac{1}{2}-\lambda}}{\sqrt{2\pi\left|\boldsymbol{\Sigma}\right|}K_{\lambda}\left(\sqrt{\chi\psi}\right)} \frac{K_{\lambda-\frac{1}{2}}\left(\sqrt{\left(\chi+\left(\boldsymbol{x}-\boldsymbol{\mu}\right)^2\boldsymbol{\Sigma}^{-1}\right)\left(\psi+\boldsymbol{\gamma}^2\boldsymbol{\Sigma}^{-1}\right)}\right)\exp\left(\left(\boldsymbol{x}-\boldsymbol{\mu}\right)\Sigma^{-1}\boldsymbol{\gamma}\right)}{\left(\sqrt{\left(\chi+\left(\boldsymbol{x}-\boldsymbol{\mu}\right)^2\boldsymbol{\Sigma}^{-1}\right)\left(\psi+\boldsymbol{\gamma}^2\boldsymbol{\Sigma}^{-1}\right)}\right)^{\frac{1}{2}-\lambda}}, \tag{3}$$

where K_v is the modified Bessel function of the third kind with index v, $\lambda \in \mathbb{R}$, $\chi, \psi \in \mathbb{R}^+$, $\boldsymbol{\mu}, \boldsymbol{\gamma} \in \mathbb{R}^N$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{N\times N}$. To gain some intuition of the expression above, one often writes the generalised hyperbolic distribution as the following mean-variance mixture. A random variable $\boldsymbol{X}$ is said to have a GH distribution if

$$\boldsymbol{X} = \boldsymbol{\mu} + \boldsymbol{\gamma}W + \sqrt{\boldsymbol{\Sigma}}\sqrt{W}\boldsymbol{Z}, \tag{4}$$

where $\boldsymbol{Z}$ is a N-dimensional normal random variable with zero mean and unit variance, and W has a generalised inverse Gaussian distribution with parameters λ, χ and ψ, i.e. $W \sim GIG(\lambda, \chi, \psi)$. Since $\boldsymbol{X}$ given $W = w$ is Normal with conditional mean $\boldsymbol{\mu}+\boldsymbol{\gamma}w$ and variance $\boldsymbol{\Sigma}w$, it is clear that $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are location and dispersion parameters, respectively. There is a further scale parameter χ, a skewness parameter ψ to allow for flexible tail modelling; and the scalar λ, which characterises certain subclasses, also influences the size of mass contained in the tails (Barndorff-Nielsen, 1977).

One of the appealing properties of normal mixtures is that the moment generating function of a GH random variable $\boldsymbol{X}$ can be easily calculated using the moment generating function of the generalised inverse Gaussian distribution. In particular, the mean and variance are given by

$$E(\boldsymbol{X}) = \boldsymbol{\mu} + \boldsymbol{\gamma}E(W) \quad \text{and} \quad Var(\boldsymbol{X}) = \boldsymbol{\gamma}^2 Var(W) + \boldsymbol{\Sigma}E(W). \tag{5}$$

The GH distribution, as the name suggests, is of a very general form, and contains as special cases many important distribution widely used in applications. It includes, among others, Student's t-distribution, the Laplace distribution, the hyperbolic distribution, the normal-inverse Gaussian distribution and the variance-gamma distribution. It is often used in economics, with particular application in the fields of modelling financial markets and risk management (Eberlein and Hammerstein, 2004), due to its flexiblity and heavy tails.

There are several parameterisations of the GH distribution. The $(\lambda, \chi, \psi, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\gamma})$ parameterisation used in this paper has the drawback of an identification problem (Barndorff-Nielsen and Shephard, 2001), i.e. $GH(\lambda, \chi, \psi, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\gamma})$ and $GH(\lambda, \chi/k, k\psi, \boldsymbol{\mu}, k\boldsymbol{\Sigma}, k\boldsymbol{\gamma})$ are identical for any $k > 0$. This is because that χ and $|\boldsymbol{\Sigma}|$ are not separately identified. This problem can be solved by introducing a suitable constraint on the parameters. Barndorff-Nielsen and Shephard (2001) fixed the determinant of $\boldsymbol{\Sigma}$ to be 1. However under this setup, it is difficult to standardise the GH distribution so that it has mean zero and unit variance. Such standardisation is necessary when developing our hierarchical model. Fortunately, the identification problem can also be solved by fixing $\chi = 1$, and that leads to a simple standardisation method (Mencia and Sentana, 2004), which we restate here in the following Proposition.

Proposition 1 *Let us modify* (3) *by replacing* $\boldsymbol{\gamma} = \boldsymbol{\Sigma\tau}$, $\boldsymbol{\tau} > 0$, *we have* $\boldsymbol{X} \sim GH(\lambda, \chi, \psi, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\tau})$. *For any* $\lambda, \psi, \boldsymbol{\tau}$ *and fixed* $\chi = 1$, *if* $\boldsymbol{\mu}$ *and* $\boldsymbol{\Sigma}$ *satisfy*

$$\boldsymbol{\mu} = -L\left(\boldsymbol{\tau}, \psi, \lambda\right)\boldsymbol{\tau} \quad \text{and} \quad \boldsymbol{\Sigma} = \frac{\sqrt{\psi}}{M_{\lambda+1}\left(\sqrt{\psi}\right)}\left(\boldsymbol{I}_N + \frac{L\left(\boldsymbol{\tau}, \psi, \lambda\right) - 1}{\boldsymbol{\tau'\tau}}\boldsymbol{\tau\tau'}\right), \tag{6}$$

where

$$M_{\lambda+1}\left(\sqrt{\psi}\right) = \frac{K_{\lambda+1}\left(\sqrt{\psi}\right)}{K_{\lambda}\left(\sqrt{\psi}\right)}, \; N_{\lambda+2}\left(\sqrt{\psi}\right) = \frac{K_{\lambda+2}\left(\sqrt{\psi}\right)K_{\lambda}\left(\sqrt{\psi}\right)}{K_{\lambda+1}\left(\sqrt{\psi}\right)^2} \quad \text{and} \tag{7}$$

$$L\left(\boldsymbol{\tau}, \psi, \lambda\right) = \frac{\sqrt{1 + 4\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)\boldsymbol{\tau'\tau}} - 1}{2\boldsymbol{\tau'\tau}\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)}, \tag{8}$$

then $E(\boldsymbol{X}) = 0$ *and* $Var(\boldsymbol{X}) = 1$.

An asymptotic justification for the use of the GH distribution can be found in Liu et al. (2014). The authors showed that for any chosen threshold, the GH distribution can approximate the power law decay of the GP distribution in the tails, if the given GP distribution has $\kappa > 0$ or $\kappa \to 0$. They also demonstrated that the GH model can achieve a good fit at the centre and shoulders of the distribution by applying the model to various rainfall datasets. In addition, fitting a GH model does not involve choosing an appropriate threshold which can be a major drawback of GP models because they do not account for the uncertainty associated with the choice of threshold, and potentially useful information below the threshold is discarded for estimating the extremes. This is a significant advantage of modelling with the GH distribution as it allows us to make inferences of not only the rainfall extremes, but also the entire distribution. Consequently the prediction credible intervals (CI) for return periods of extreme events and other statistical inference obtained under the GH model have a larger sample contributing to the estimates.

1.4 Paper outline

The next section describes the precipitation data sources and some preliminary analysis. In Section 3 we describe the structure of our Bayesian hierarchical model. We discuss the GH-based hierarchical model for rainfall and rainfall extremes in Section 3.1, discuss the Bayesian framework including prior distributions in Section 3.2, briefly describe our MCMC method for model inference in Section 3.3, and discuss how our model was used to produce spatial predictions in Section 3.4, and briefly illustrate temporal prediction in Section 3.5. We then present results, including spatial predictions and model validation in Section 4. Finally, we conclude with a discussion in Section 5.

2 Data

Winter rainfall in south-west Western Australia (SWWA) was once considered the most consistent and reliable anywhere in Australia. However, around the mid 1970s, there was a shift to more volatile winter conditions, which has continued to this day, so the region can experience both extreme and sometimes insufficient rainfall in winter (Ruprecht et al., 2005). These conditions have had strong negative impacts on urban infrastructure, surface and ground water supplies, agriculture and natural ecosystems.

Although these events are rare, understanding the frequency, intensity and duration is important for public safety and long-term planning. In this paper we illustrate the proposed method by applying it to precipitation data from 21 weather stations in the Natural Resource Management (NRM) Region 504 of

SWWA; see Figure 1. Our method aims to produce out-of-sample spatial predictions of precipitation return periods, duration over/below thresholds and relevant prediction maps.

2.1 Study region, weather stations, and covariates

Daily precipitation (mm) data are available from 21 meteorological monitoring sites in the study region obtained from the Australian Bureau of Meteorology website (http://www.bom.gov.au); see Figure 1. In this paper, we analyse daily precipitation data for 92 days in the months of June-August for 48 years from 1960 to 2008. Hence, a total of 92,736 data points for daily precipitation are used in this analysis, of which 1.10% were missing, and 38.27% were recorded 0 mm.

This paper also accounts for various climatic drivers that have an influence on precipitation levels.

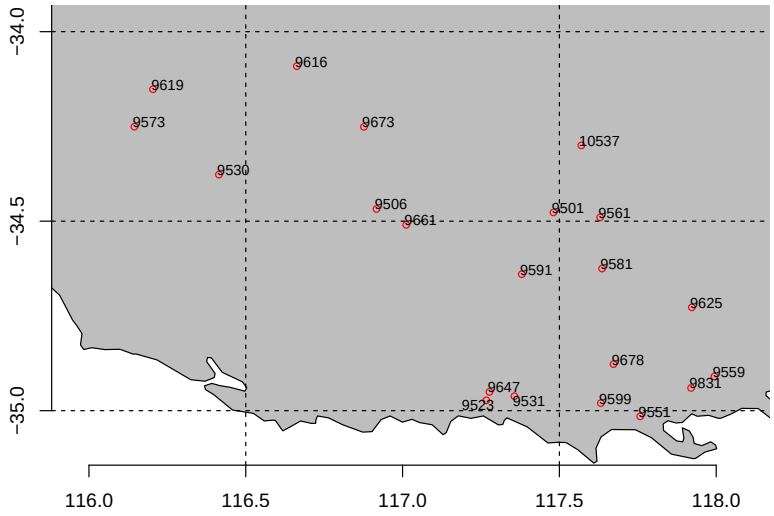

Figure 1: Weather stations in the the south-western agricultural zone of Western Australia (NRM Region 504). Location numbers are given next to the circles.

Understanding the effects of climatic drives on rainfall and the pattern of seasonal precipitation is important for assessing agricultural productivity and water risk-management. Empirical evidences on the relationships between them are discussed in many studies, e.g. Schepen et al. (2012). Three important climatic drivers: the El Niño southern oscillation anomaly (NINO 3.4), Southern Hemisphere Annular Mode (SAMI) and the Indian Ocean Dipole (IOD) are considered in this study to understand their effects on precipitation levels. It has been observed that these indices have potential teleconnection with precipitation levels. The NINO 3.4 index is a monthly time series of mean sea surface temperature (SST) from an equatorial region of the eastern Pacific that covers 5° south to 5° north and 125° west to 175° east (Schepen et al., 2012). The SAMI index includes the Antarctic oscillation and is obtained by the differences in the normalised monthly zonal mean sea level pressure between 40° and 70° south (Nan and Li, 2003). The IOD index is calculated from monthly SST anomalies in the equatorial Indian Ocean (Saji et al., 1999). Another important variable, elevation (ELE), has also been considered in our model as it has a strong correlation with rainfall in many places.

We have not accounted for any seasonal effects in our data. Restricting our analysis to the winter months reduces seasonality and on inspecting the data from several sites there was no obvious seasonal effect. Likewise, we have not accounted for any temporal trends in the data. However, our series are relatively short and it may be difficult to discover any trend in winter precipitation over the last 45 years. Furthermore, the purpose of this study is not focused on climate change. Both climate change and seasonal

effects would be simple extensions of this study. For example, the model proposed in this paper can be generalised to a dynamic linear model (West et al., 1985; Stroud et al., 2001), which is popular for modelling data with seasonal variations. Recent applications include Dou et al. (2010); Ghosh et al. (2010); Mahmoudian and Mohammadzadeh (2014) and Sahu and Bakar (2012).

2.2 Exploratory analysis

Table 1 provides summary statistics for the variables used in the model. Daily rainfall ranges from 0 mm to 112.00 mm with a mean value of 3.73 mm. The table also provides summary statistics for the explanatory variable. NINO 3.4 ranges from 24.43 to 29.14, the SAMI ranges from -7.13 to 5.36 and the IOD varies from -2.74 to 3.55 units for the 45-year period considered in this study. Finally the highest elevation of 300 metre is at location 3 in Figure 1.

The wettest year was 1987 with an average winter daily rainfall of 4.18 mm, while the dryest was 1974,

Variables	Minimum	Mean	Maximum
Rainfall (mm)	0.00	3.73	112.00
NINO 3.4	24.43	27.01	29.14
SAMI	-7.13	0.14	5.36
IOD	-2.74	0.01	3.55
ELE (metre)	10.00	185.40	300.00

Table 1: Summary statistics for the variables used.

with an average winter daily rainfall of 1.12 mm. The largest observed number of days with zero rainfall (33 days) occurred in 1978 and the smallest (17 days) in 1986 and 2005. The data shows temporal dependencies and non-stationarity. Also the level of rainfall differs from site to site which may lead to heavier rainfall or drier conditions in some areas.

3 Model specification

In this section we describe the general structure of the proposed hierarchical model in detail. Hierarchical models have been widely used in spatio-temporal modelling and they allow one to statistically model a complex process and its relationship to observations in several simple components. For an introduction to such models, see Gelman et al. (2013).

3.1 General structure of the hierarchical model

Let $D_t(s_i)$ denote the observed point referenced data at location s_i, $i = 1, \ldots, N$, and time t, $t = 1, \ldots, N$. We are interested in making inference on the D process on the basis of data. There are three layers in our hierarchical model. In the first level of modelling we write:

$$D_t(s_i) = Y_t(s_i) I_{\{Y_t(s_i) \in E\}}, \tag{9}$$

where $I_{\{\cdot\}}$ is the standard indicator function, and the set $E = (\infty, 0]$ contains values for which $Y_t(s_i)$ is observed. Similiar censoring approach for modelling zero precipitation was also used in Stein (1992), Glasbey and Nevison (1997) and Yuen and Guttorp (2014). We will discuss alternative methods in Section

5. The mean process is modelled in the second level of the hierarchy:

$$Y_t(s_i) = \beta_0 + \sum_{k=1}^{M} \beta_k X_{kt}(s_i) + V_t(s_i), \tag{10}$$

where β_0 represents the global intercept, $X_{kt}(s_i)$ is the kth regressor with regression coefficient β_k. In the final stage of modelling we have the process for the spatially and temporally correlated error $V_t(s_i)$:

$$V_t(s_i) = bV_{t-1}(s_i) + U_t(s_i), \tag{11}$$

$$U_t(s_i) = \sum_{j=1}^{N} \sigma_t(s_i, s_j) Z_t(s_j), \tag{12}$$

$$\sigma_t^2(s_i, s_i) = \omega + \alpha U_{t-1}^2(s_i), \quad \text{with } \omega, \alpha \geq 0, \tag{13}$$

where $Z_t(s_j)$ is the jth component of a N-dimensional standardised generalised hyperbolic distribution, $\sigma_t^2(s_i, s_j)$ (fully defined in Section 3.1.2) is the covariance between s_i and s_j, and $\sigma_t^2(s_i, s_i)$ is the spatially and temporally varying variance process. For simplicity, the autoregressive parameter is modelled as a single global effect b, whereas a more general structure can be introduced and is discussed in Section 5.

Writing the above in vector and matrix notations, let $\boldsymbol{Y}_t = \{Y_t(s) : s \in R \subset \mathbb{R}^N\}$ be a real-valued spatial process, we have the following spatio-temporal model, for $t = 1, \ldots, T$:

$$\boldsymbol{D}_t = \boldsymbol{A}\boldsymbol{Y}_t, \tag{14}$$

$$\boldsymbol{Y}_t = \beta_0 + \sum_{k=1}^{M} \boldsymbol{X}_{kt}\beta_k + \boldsymbol{V}_t, \tag{15}$$

$$\boldsymbol{V}_t = b\boldsymbol{V}_{t-1} + \boldsymbol{U}_t, \tag{16}$$

$$\boldsymbol{U}_t = \boldsymbol{\sigma}_t \boldsymbol{Z}_t, \tag{17}$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{S}_t \boldsymbol{\rho} \boldsymbol{S}_t, \tag{18}$$

where the matrices and vectors are defined in the following subsections.

3.1.1 The mean process

The mean process is modelled by (14) and (15). For a finite number of spatial locations $\{s_1, \ldots, s_N\} \subset R$, the vector $\boldsymbol{D}_t = (D_t(s_1), \ldots, D_t(s_N))'$ is the observed precipitation at weather stations at time t, $\boldsymbol{Y}_t = (Y_t(s_1), \ldots, Y_t(s_N))'$ represents the latent underlying process at time t, $\boldsymbol{X}_{kt} = (X_{kt}(s_1), \ldots, X_{kt}(s_N))'$ for $k = 1, \ldots, M$ are the vector processes of M known covariates at time t, and $\boldsymbol{V}_t = (V_t(s_1), \ldots V_t(s_N))'$ are the spatially correlated errors. We note that $\boldsymbol{D}_t, \boldsymbol{Y}_t, \boldsymbol{X}_{kt}, \boldsymbol{V}_t \in \mathbb{R}^N$ and that the $N \times N$ matrix $\boldsymbol{A}$ is a diagonal matrix whose diagonal entries starting in the upper left corner are given by standard indicator functions $I_{\{Y_t(s_i) \in E\}}, \ldots, I_{\{Y_t(s_N) \in E\}}$. As we have already discussed, for simplicity $\beta_0 \in \mathbb{R}$, $\beta_k \in \mathbb{R}$ are the global intercept and coefficient for the kth covariate respectively.

We assume the spatial correlation matrix $\boldsymbol{\rho}$ has Matérn structure (Matérn, 1986), independent of time, i.e. the (i, j) component of the matrix $\boldsymbol{\rho}$ is

$$\rho(s_i, s_j; \theta, \phi) = \frac{1}{2^{\phi-1}\Gamma(\phi)} \left(2\sqrt{\phi}\, \|s_i - s_j\|\, \theta\right)^{\phi} K_{\phi}\left(2\sqrt{\phi}\, \|s_i - s_j\|\, \theta\right), \quad \theta, \phi > 0, \tag{19}$$

where $\Gamma(\phi)$ is the standard gamma function, K_ϕ is the modified Bessel function of second kind with order θ, and $\|s_i - s_j\|$ is the distance between s_i and s_j. The parameter θ controls the rate of decay of the correlation as the distance $\|s_i - s_j\|$ increases and the parameter ϕ controls smoothness of the random field. Our covariance model assmes the process is isotropic and stationary; we found it impossible to detect any nonstationary or anisotropy with only 21 stations. For convenience we define one common variance σ_β^2 for all spatially varying parameters, i.e. $\sigma_k = \sigma_\beta$ for $k = 1, \ldots, M$.

3.1.2 Spatially and temporally correlated errors

The error process $\boldsymbol{V}_t$ is modelled by (16), (17) and (18), where $\boldsymbol{Z}_t \in \mathbb{R}^N$ is a standardised multivariate generalised hyperbolic random variable and $\boldsymbol{\sigma}_t$ is a lower trianglar $N \times N$ matrix given by the Cholesky decomposition of the $N \times N$ covariance process matrix $\boldsymbol{\sigma}_t^2$. In its covariance structure given in (18), the variance matrix $\boldsymbol{S}_t$ is a $N \times N$ diagonal matrix whose diagonal entries starting in the upper left corner are given by the ARCH process $\sqrt{\omega + \alpha \boldsymbol{U}_{t-1}^2}$ with $\omega, \alpha \in \mathbb{R}$, whereas the correlation matrix is given by $\boldsymbol{\rho}_t \in \mathbb{R}^{N \times N}$. In detail, we write

$$
\begin{aligned}
\boldsymbol{\sigma}_t^2 =& \boldsymbol{S}_t \boldsymbol{\rho} \boldsymbol{S}_t \quad \text{with} \\
=& \begin{pmatrix} \sigma_t(s_1) & 0 & \cdots & 0 \\ 0 & \sigma_t(s_2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_t(s_N) \end{pmatrix} \begin{pmatrix} 1 & \rho(s_1, s_2) & \cdots & \rho(s_1, s_N) \\ \rho(s_2, s_1) & 1 & \cdots & \rho(s_2, s_N) \\ \vdots & \vdots & \ddots & \vdots \\ \rho(s_N, s_1) & \rho(s_N, s_2) & \cdots & 1 \end{pmatrix} \\
& \begin{pmatrix} \sigma_t(s_1) & 0 & \cdots & 0 \\ 0 & \sigma_t(s_2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_t(s_N) \end{pmatrix},
\end{aligned} \tag{20}
$$

where the spatial correlation matrix is given by (19). If we write

$$
\text{Chol} \begin{pmatrix} 1 & \rho(s_1, s_2) & \cdots & \rho(s_1, s_N) \\ \rho(s_2, s_1) & 1 & \cdots & \rho(s_2, s_N) \\ \vdots & \vdots & \ddots & \vdots \\ \rho(s_N, s_1) & \rho(s_N, s_2) & \cdots & 1 \end{pmatrix} = \begin{pmatrix} l_{11} & 0 & \cdots & 0 \\ l_{21} & l_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ l_{N1} & l_{N2} & \cdots & l_{NN}, \end{pmatrix} = \boldsymbol{L} \tag{21}
$$

where $\boldsymbol{L}$ is a lower triangular matrix with

$$
l_{kk} = \sqrt{1 - \sum_{j=1}^{k-1} l_{kj}^2} \quad \text{and} \quad l_{ik} = \frac{1}{l_{kk}} \left(\rho(s_i, s_k) - \sum_{j=1}^{k-1} l_{ij} l_{kj} \right), \tag{22}
$$

we have

$$
\boldsymbol{\sigma}_t = \boldsymbol{S}_t \boldsymbol{L}. \tag{23}
$$

3.1.3 Standardised generalised hyperbolic random variables

The random vector $\boldsymbol{Z}_t \in \mathbb{R}^N$ follows a standardiased multivariate generalised hyperbolic distribution with parameters $\lambda, \psi, \chi \in \mathbb{R}$, $\boldsymbol{\mu}, \boldsymbol{\tau} \in \mathbb{R}^N$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$. That is $\boldsymbol{Z}_t \sim MVGH\left(\lambda, \psi, \chi, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\tau}\right)$ with

$$\chi = 1, \quad \boldsymbol{\mu} = -L\left(\boldsymbol{\tau}, \psi, \lambda\right)\boldsymbol{\tau} \quad \text{and} \quad \boldsymbol{\Sigma} = \frac{\sqrt{\psi}}{M_{\lambda+1}\left(\sqrt{\psi}\right)}\left(\boldsymbol{I}_N + \frac{L\left(\boldsymbol{\tau}, \psi, \lambda\right) - 1}{\boldsymbol{\tau}'\boldsymbol{\tau}}\boldsymbol{\tau}\boldsymbol{\tau}'\right), \tag{24}$$

where $M_{\lambda+1}\left(\sqrt{\psi}\right)$ and $L\left(\boldsymbol{\tau}, \psi, \lambda\right)$ are given by (7) and (8). Note that $\boldsymbol{\tau}$ is the $N \times 1$ skewness parameter. In this paper we assume that $\tau(s_i) = m_\tau$ for $i = 1, \ldots, N$. We note that a spatial random effect on $\boldsymbol{\tau}$ can be introduced. However this might lead to a non-identifiability problem because there will be two sets of random effects on the overall skewness, including the covariance process $\boldsymbol{\sigma}_t^2$, only the product of which is identified by the data; see (28). Routine calculation then yields the conditional distribution for $\boldsymbol{Y}_t$:

$$\boldsymbol{Y}_t | \boldsymbol{\zeta}, \boldsymbol{Y}_{t-1} \sim MVGH\left(\lambda, \psi, 1, \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t, \boldsymbol{\tau}_t\right), \tag{25}$$

where $\boldsymbol{\zeta} = (a, b, \alpha, \omega, \lambda, \psi, \boldsymbol{\tau}, \beta_1, \ldots, \beta_k, \theta)$ represents all parameters in the model and

$$\boldsymbol{\mu}_t = \beta_0 + \sum_{k=1}^{M} \boldsymbol{X}_{kt}\boldsymbol{\beta}_k + b\boldsymbol{V}_{t-1} + \boldsymbol{\sigma}_t\boldsymbol{\mu}, \tag{26}$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\sigma}_t\boldsymbol{\Sigma}\boldsymbol{\sigma}_t', \tag{27}$$

$$\boldsymbol{\tau}_t = \boldsymbol{\sigma}_t\boldsymbol{\tau}. \tag{28}$$

3.2 Bayesian framework

Inference for the parameters $\boldsymbol{\zeta}$ in our model given the underlying spatio-temporal process $\boldsymbol{Y} = (\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_T)$ comes simply from Bayes rule:

$$p(\boldsymbol{\zeta}|\boldsymbol{Y}) \propto p(\boldsymbol{Y}|\boldsymbol{\zeta})p(\boldsymbol{\zeta}), \tag{29}$$

where p denotes a probability density.

Since at any given time t, some components of $\boldsymbol{Y}_t$ may be censored, to compute the likelihood, we require a data argumentation procedure to recover censored $\boldsymbol{Y}_t$ components. To be precise, let $\boldsymbol{Y}_{O_t}$ where $O_t \subset \{1, \ldots, N\}$ indicate the subsets of the process at time t that occur above zero, whereas $\boldsymbol{Y}_{C_t}$ with $C_t = \{1, \ldots, N\} \backslash O_t$ are the censored part of $\boldsymbol{Y}_t$, then the observed information $\mathscr{I}_t$ at time t is

$$\mathscr{I}_t = \{Y_t(s_i), i \in O_t\} \cup \{Y_t(s_i) < 0, i \in C_t\}. \tag{30}$$

We note that $\boldsymbol{D}_t$ carries similar information regarding $\boldsymbol{Y}_t$ as $\mathscr{I}_t$, but we observe 0 for $i \in C_t$ rather than only knowing $Y_t(s_i) < 0$. The likelihood for $\boldsymbol{\zeta} = (a, b, \alpha, \omega, \lambda, \psi, m_\tau, \beta_0, \boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_k, \theta)$ given $\mathscr{I}_t$ and $\boldsymbol{Y}_{t-1}$ is

$$L(\boldsymbol{\zeta}|\mathscr{I}_t, \boldsymbol{Y}_{t-1}) = p(\boldsymbol{Y}_{O_t}|\boldsymbol{\zeta}, \boldsymbol{Y}_{t-1}) \int_{\boldsymbol{y}_t \leq \boldsymbol{0}} p(\boldsymbol{y}_t|\boldsymbol{Y}_{O_t}, \boldsymbol{\zeta}, \boldsymbol{Y}_{t-1}) d\boldsymbol{y}_t, \tag{31}$$

which can be computed using a data augmentation method by embedding Monte Carlo integration within our Metropolis-Hastings algorithm. A similar method was also used in Liu et al. (2014) and De Oliveira (2005).

3.2.1 Prior distributions

We assign priors to the model parameter $\boldsymbol{\zeta}$. We do not assume any prior knowledge on how the covariates are related to precipitation, and thus we choose uninformative priors for the parameters $\beta_1, \ldots, \beta_M$ by considering $\beta_k \sim N(m, v)$, for all k, with zero mean and very large variance. In this application, we rely on knowledge of the space in which we model to set priors for the spatial decay parameter θ (Banerjee et al., 2004; Cooley et al., 2007). Since we model in the latitude/longitude space, we use Unif$(0, 1.5)$ as our prior, which sets the maximum range of the exponential variogram model to be approximately 200 km.

We also assign uninformative priors to the parameters of the generalised hyperbolic processes similar to the priors for β_k. A sensitivity analysis has also been done using different hyper-parameter values of the prior distributions; see details in Section 4.2.

3.2.2 Sampling from the posterior

For each $t \in \{1, \ldots, T\}$, the hierarchical model described above yields the following posterior density

$$
\begin{aligned}
p(\boldsymbol{Y}_{C_t}, \boldsymbol{\zeta} | \mathscr{I}_t, \boldsymbol{Y}_{t-1}) \propto & p\left(\mathscr{I}_t, \boldsymbol{Y}_{C_t}, \boldsymbol{\zeta} | \boldsymbol{Y}_{t-1}\right) \propto p(\mathscr{I}_t | \boldsymbol{Y}_{C_t}, \boldsymbol{\zeta}, \boldsymbol{Y}_{t-1}) p(\boldsymbol{Y}_{C_t} | \boldsymbol{\zeta}) p(\boldsymbol{\zeta}) \\
& \propto \frac{\left(\sqrt{\psi/\chi}\right)^{\lambda} \left(\psi + \boldsymbol{\gamma}_t' \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\gamma}_t\right)^{\frac{N}{2}-\lambda}}{(2\pi)^{\frac{N}{2}} |\boldsymbol{\Sigma}_t|^{\frac{1}{2}} K_{\lambda}\left(\sqrt{\psi/\chi}\right)} \\
& \quad \times \frac{K_{\lambda-\frac{N}{2}}\left(\sqrt{(\chi + Q(\boldsymbol{Y}_t))\left(\psi + \boldsymbol{\gamma}_t' \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\gamma}_t\right)}\right) \exp\left(\boldsymbol{Y}_t - \boldsymbol{\mu}_t\right)' \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\gamma}_t}{\left(\sqrt{(\chi + Q(\boldsymbol{Y}_t))\left(\psi + \boldsymbol{\gamma}_t' \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\gamma}_t\right)}\right)^{\frac{N}{2}-\lambda}} \\
& \quad \times \prod_{i \in C_t} \mathbf{1}_{\{Y_t(s_i) < 0\}} \times \pi(b) \pi(\beta_0) \prod_{k=1}^{M} \pi(\beta_k) \pi(\omega) \pi(\alpha) \pi(\lambda, \psi) \pi(m_\tau) \pi(\theta) \pi(\phi), \quad (32)
\end{aligned}
$$

where $Q(\boldsymbol{Y}_t) = (\boldsymbol{Y}_t - \boldsymbol{\mu}_t)' \boldsymbol{\Sigma}_t^{-1} (\boldsymbol{Y}_t - \boldsymbol{\mu}_t)$, $\boldsymbol{\gamma}_t = \boldsymbol{\Sigma}_t \boldsymbol{\tau}_t$, and $\boldsymbol{\mu}_t$, $\boldsymbol{\Sigma}_t$ and $\boldsymbol{\tau}_t$ are given by (26), (27) and (28). We note that $\boldsymbol{Y}_t = (\boldsymbol{Y}_{O_t}, \boldsymbol{Y}_{C_t})$ is the underlying vector, whose censored components $\{Y_t(s_i), i \in C_t\}$ are initialised at the censoring limit zero and then sampled collectively at each iteration of the Metropolis-Hastings algorithm from their full conditional distribution. Here we assume that $\boldsymbol{Y}_{C_{t-1}}$ has been recovered, so $\boldsymbol{Y}_{t-1}$ is avaliable.

Let $\boldsymbol{\zeta}^{[j]} = \left(a^{[j]}, b^{[j]}, \alpha^{[j]}, \omega^{[j]}, \lambda^{[j]}, \psi^{[j]}, m_\tau^{[j]}, \beta_0^{[j]}, \beta_1^{[j]}, \ldots, \beta_k^{[j]}, \theta^{[j]}, \phi^{[j]}\right)$ be the jth realisation in the parameter space. We can rewrite (25) as $\boldsymbol{Y}_t | \boldsymbol{\zeta}^{[j]} = (\boldsymbol{Y}_{O_t}, \boldsymbol{Y}_{C_t}) | \boldsymbol{\zeta}^{[j]} \sim MGH\left(\lambda^{[j]}, \psi^{[j]}, 1, \boldsymbol{\mu}_t^{[j]}, \boldsymbol{\Sigma}_t^{[j]}, \boldsymbol{\tau}_t^{[j]}\right)$, where $\boldsymbol{\mu}_t^{[j]}$, $\boldsymbol{\tau}_t^{[j]}$ and $\boldsymbol{\Sigma}_t^{[j]}$ are given by (26), (27) and (28) with parameters $\boldsymbol{\zeta}^{[j]}$. If we write

$$
\boldsymbol{\mu}_t^{[j]} = \begin{pmatrix} \boldsymbol{\mu}_{O_t} \\ \boldsymbol{\mu}_{C_t} \end{pmatrix}, \ \boldsymbol{\tau}_t^{[j]} = \begin{pmatrix} \boldsymbol{\tau}_{O_t} \\ \boldsymbol{\tau}_{C_t} \end{pmatrix} \text{ and } \boldsymbol{\Sigma}_t^{[j]} = \begin{pmatrix} \boldsymbol{\Sigma}_{O_t} & \boldsymbol{\Sigma}_{OC_t} \\ \boldsymbol{\Sigma}_{COt} & \boldsymbol{\Sigma}_{C_t} \end{pmatrix}, \quad (33)
$$

then the censored components of $\boldsymbol{Y}_t$ given $\boldsymbol{\zeta}^{[j]}$ and the observed information $\mathscr{I}_t$ follow a multivariate GH distribution:

$$
\begin{aligned}
p(\boldsymbol{Y}_{C_t} | \boldsymbol{\zeta}^{[j]}, \mathscr{I}_t) = & p(\boldsymbol{Y}_{C_t} | \boldsymbol{\zeta}^{[j]}, \boldsymbol{Y}_{O_t}) \\
= & MVGH\left(\lambda', \chi', \psi', \boldsymbol{\mu}_t', \boldsymbol{\Sigma}_t', \boldsymbol{\gamma}_t'\right) \times \mathbf{1}_{\{\boldsymbol{Y}_t < \mathbf{0}\}}. \quad (34)
\end{aligned}
$$

Here

$$\lambda' = \lambda - \dim\left(\boldsymbol{Y}_{O_t}\right)/2 \tag{35}$$

$$\boldsymbol{\mu}'_t = \boldsymbol{\mu}_{C_t} + \left(\boldsymbol{Y}_{O_t} - \boldsymbol{\mu}_{O_t}\right)\boldsymbol{\Sigma}_{O_t}^{-1}\boldsymbol{\Sigma}_{OC_t} \tag{36}$$

$$\boldsymbol{\Sigma}'_t = \left|\boldsymbol{\Sigma}_{O_t}\right|^{1/\dim(\boldsymbol{Y}_{C_t})}\left(\boldsymbol{\Sigma}_{C_t} - \boldsymbol{\Sigma}_{CO_t}\boldsymbol{\Sigma}_{C_t}^{-1}\boldsymbol{\Sigma}_{OC_t}\right) \tag{37}$$

$$\boldsymbol{\gamma}'_t = \left(\boldsymbol{\gamma}_{C_t} - \boldsymbol{\gamma}_{O_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)\left(\boldsymbol{\Sigma}_{C_t} - \boldsymbol{\Sigma}_{CO_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)^{-1}\boldsymbol{\Sigma}'_t \tag{38}$$

$$\chi' = \left|\boldsymbol{\Sigma}_{O_t}\right|^{1/\dim(\boldsymbol{Y}_{C_t})}\left(1 + \left(\boldsymbol{Y}_{O_t} - \boldsymbol{\mu}_{O_t}\right)\boldsymbol{\Sigma}_{O_t}^{-1}\left(\boldsymbol{Y}_{O_t} - \boldsymbol{\mu}_{O_t}\right)'\right) \tag{39}$$

$$\psi' = \left|\boldsymbol{\Sigma}_{O_t}\right|^{1/N}\left(\psi + \boldsymbol{\gamma}'_{O_t}\boldsymbol{\Sigma}_{O_t}^{-1}\boldsymbol{\gamma}_{O_t}\right)\left|\boldsymbol{\Sigma}_{O_t}\right|^{1/\dim(\boldsymbol{Y}_{C_t})} - \tag{40}$$

$$\left(\left(\boldsymbol{\gamma}_{C_t} - \boldsymbol{\gamma}_{O_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)\left(\boldsymbol{\Sigma}_{C_t} - \boldsymbol{\Sigma}_{CO_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)^{-1}\right)'\boldsymbol{\Sigma}'_t\left(\boldsymbol{\gamma}_{C_t} - \boldsymbol{\gamma}_{O_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)\left(\boldsymbol{\Sigma}_{C_t} - \boldsymbol{\Sigma}_{CO_t}\boldsymbol{\Sigma}_{O_t}^{-1}\right)^{-1}, \tag{41}$$

with $\boldsymbol{\gamma}_{O_t} = \boldsymbol{\Sigma}_{O_t}\boldsymbol{\tau}_{O_t} + \boldsymbol{\Sigma}_{OC_t}\boldsymbol{\tau}_{C_t}$ and $\boldsymbol{\gamma}_{C_t} = \boldsymbol{\Sigma}_{CO_t}\boldsymbol{\tau}_{O_t} + \boldsymbol{\Sigma}_{C_t}\boldsymbol{\tau}_{C_t}$. See Section 3.3 for further details.

Data augmentation is done for all t in a recursive manner, so $\boldsymbol{Y}_{C_t}$ is recovered and $\boldsymbol{Y}_t$ is available for all t. This method for censored observations ensures that likelihood distributions given data $\mathscr{I}_t$ follow (31) (De Oliveira, 2005). Under the assumption that $\boldsymbol{Y}_t$ and $(\boldsymbol{Y}_{C_t}, \mathscr{I}_t)$ carry the same information, the posterior density has the form

$$\begin{aligned} p\left(\boldsymbol{Y}_{C_T}, \boldsymbol{Y}_{C_{T-1}}, \ldots, \boldsymbol{Y}_{C_1}, \boldsymbol{\zeta} | \mathscr{I}_T, \mathscr{I}_{T-1}, \ldots, \mathscr{I}_1\right) &= \prod_{t=1}^{T} p\left(\boldsymbol{Y}_{C_t}, \boldsymbol{\zeta} | \mathscr{I}_t, \boldsymbol{Y}_{C_{t-1}}, \mathscr{I}_{t-1}\right) \times p\left(\boldsymbol{Y}_{C_0}, \mathscr{I}_0\right) \\ &= \prod_{t=1}^{T} p\left(\boldsymbol{Y}_{C_t}, \boldsymbol{\zeta} | \mathscr{I}_t, \boldsymbol{Y}_{t-1}\right) \times p\left(\boldsymbol{Y}_0\right). \end{aligned} \tag{42}$$

Due to the recursive nature of the hierarchical model, full conditionals are non-standard and require Metropolis-Hastings sampling. See details in Gelman et al. (2013) and Tierney (1994).

3.3 Spatial predictive distribution

Recall that our primary goal is to make spatial predictions for out-of-sample locations given the observed data. In practice posterior predictive distributions can be produced by sampling from a finite set of locations $\boldsymbol{s}^*$, conditional on $\boldsymbol{Y}$ and $\boldsymbol{\zeta}$. We want to predict $\boldsymbol{Y}_t^* = (Y_t(s_1^*), \ldots, Y_t(s_r^*))$ at locations $\boldsymbol{s}^* = \{s_1^*, \ldots, s_r^*\}$ at time t. Let us consider the following hierarchical model:

$$\widetilde{\boldsymbol{Y}}_t = \beta_0 + \sum_{k=1}^{M} \widetilde{\boldsymbol{X}}_{kt}\beta_k + \widetilde{\boldsymbol{V}}_t, \tag{43}$$

$$\widetilde{\boldsymbol{V}}_t = b\widetilde{\boldsymbol{V}}_{t-1} + \widetilde{\boldsymbol{U}}_t, \tag{44}$$

$$\widetilde{\boldsymbol{U}}_t = \widetilde{\boldsymbol{\sigma}}_t\widetilde{\boldsymbol{Z}}_t, \tag{45}$$

$$\widetilde{\boldsymbol{\sigma}}_t^2 = \widetilde{\boldsymbol{S}}\widetilde{\boldsymbol{\rho}}_t\widetilde{\boldsymbol{S}}, \tag{46}$$

where

$$\widetilde{\boldsymbol{Y}}_t = \begin{pmatrix} \boldsymbol{Y}_t \\ \boldsymbol{Y}_t^* \end{pmatrix}, \widetilde{\boldsymbol{X}}_{kt} = \begin{pmatrix} \boldsymbol{X}_{kt} \\ \boldsymbol{X}_{kt}^* \end{pmatrix}, \widetilde{\boldsymbol{V}}_t = \begin{pmatrix} \boldsymbol{V}_t \\ \boldsymbol{V}_t^* \end{pmatrix}, \widetilde{\boldsymbol{U}}_t = \begin{pmatrix} \boldsymbol{U}_t \\ \boldsymbol{U}_t^* \end{pmatrix}, \text{ and } \widetilde{\boldsymbol{Z}}_t = \begin{pmatrix} \boldsymbol{Z}_t \\ \boldsymbol{Z}_t^* \end{pmatrix}, \tag{47}$$

where $\widetilde{\boldsymbol{Z}}_t \sim MGH\left(\lambda, \psi, \chi, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}, \widetilde{\boldsymbol{\tau}}\right)$ satisfying the standardisation condition. The generalised hyperbolic parameters λ, χ and ψ are the same as before, and $\widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}$ and $\widetilde{\boldsymbol{\tau}}$ are given by (24). We have $\widetilde{\boldsymbol{Y}}_t, \widetilde{\boldsymbol{X}}_t, \widetilde{\boldsymbol{V}}_t, \widetilde{\boldsymbol{U}}_t, \widetilde{\boldsymbol{Z}}_t \in \mathbb{R}^{N+r}$ and $b, \beta_0, \beta_k \in \mathbb{R}$. The conditional variance $\widetilde{\boldsymbol{\sigma}}_t^2$ of $\widetilde{\boldsymbol{U}}_t$ becomes

$$\widetilde{\boldsymbol{\sigma}}_t^2 = \begin{pmatrix} \boldsymbol{\sigma}_t^2 & \boldsymbol{\sigma}_t(\boldsymbol{s}^*, \boldsymbol{s}) \\ \boldsymbol{\sigma}_t(\boldsymbol{s}, \boldsymbol{s}^*) & \boldsymbol{\sigma}_t^{*2} \end{pmatrix}, \tag{48}$$

and consequently,

$$\widetilde{\boldsymbol{\sigma}}_t = \widetilde{\boldsymbol{S}}\,\mathrm{Chol}\left(\widetilde{\boldsymbol{\rho}}\right) = \begin{pmatrix} \boldsymbol{\sigma}_t\left(\boldsymbol{s}\right) & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\sigma}_t\left(\boldsymbol{s}^*\right) \end{pmatrix} \mathrm{Chol} \begin{pmatrix} \boldsymbol{\rho} & \boldsymbol{\rho}(\boldsymbol{s}^*, \boldsymbol{s}) \\ \boldsymbol{\rho}(\boldsymbol{s}, \boldsymbol{s}^*) & \boldsymbol{\rho}^*, \end{pmatrix} \tag{49}$$

where $\boldsymbol{\sigma}_t\left(\boldsymbol{s}^*\right) \in \mathbb{R}^{r\times r}$ is a diagonal matrix whose entries are given by $\sqrt{\omega + \alpha \boldsymbol{U}_{t-1}^{*2}}$, and $\boldsymbol{\rho}(\boldsymbol{s}^*, \boldsymbol{s}) \in \mathbb{R}^{N\times r}$, $\boldsymbol{\rho}(\boldsymbol{s}, \boldsymbol{s}^*) \in \mathbb{R}^{r\times N}$ and $\boldsymbol{\rho}^* \in \mathbb{R}^{r\times r}$ are given by (19). Hence we have

$$\widetilde{\boldsymbol{Y}}_t | \boldsymbol{\zeta}^\dagger, \widetilde{\boldsymbol{Y}}_{t-1}, = \begin{pmatrix} \boldsymbol{Y}_t \\ \boldsymbol{Y}_t^* \end{pmatrix} | \begin{pmatrix} \boldsymbol{Y}_{t-1} \\ \boldsymbol{Y}_{t-1}^* \end{pmatrix} \sim MGH\left(\lambda, \chi, \psi, \boldsymbol{\mu}_t^\dagger, \boldsymbol{\Sigma}_t^\dagger, \boldsymbol{\gamma}_t^\dagger\right), \tag{50}$$

where $\boldsymbol{\zeta}^\dagger$ represents all parameters in the model, $\boldsymbol{\gamma}_t^\dagger = \boldsymbol{\Sigma}_t^\dagger \boldsymbol{\tau}_t^\dagger = (\gamma_{1t}\ \ \boldsymbol{\gamma}_{0t})'$, and

$$\boldsymbol{\mu}_t^\dagger = \beta_0 + \sum_{k=1}^{M} \widetilde{\boldsymbol{X}}_{kt}\beta_k + b\widetilde{\boldsymbol{V}}_{t-1} + \widetilde{\boldsymbol{\sigma}}_t\widetilde{\boldsymbol{\mu}} = \begin{pmatrix} \boldsymbol{\mu}_{1t} \\ \boldsymbol{\mu}_{0t} \end{pmatrix},\ \boldsymbol{\tau}_t^\dagger = \widetilde{\boldsymbol{\sigma}}_t\widetilde{\boldsymbol{\tau}} = \begin{pmatrix} \boldsymbol{\tau}_{1t} \\ \boldsymbol{\tau}_{0t} \end{pmatrix} \text{ and } \boldsymbol{\Sigma}_t^\dagger = \widetilde{\boldsymbol{\sigma}}_t\widetilde{\boldsymbol{\Sigma}}\widetilde{\boldsymbol{\sigma}}_t' = \begin{pmatrix} \boldsymbol{\Sigma}_{1t} & \boldsymbol{\Sigma}_{10t} \\ \boldsymbol{\Sigma}_{01t} & \boldsymbol{\Sigma}_{0t}, \end{pmatrix}.$$

Due to the recursive temporal structure of the hierarchical model, we assume $\boldsymbol{Y}_{t-1}^*$ is sampled in previous step and is known. We also assume that censored components of $\boldsymbol{Y}_t$ and $\boldsymbol{Y}_{t-1}$ are recovered using (34). Here we have $\boldsymbol{\mu}_{1t}, \boldsymbol{\gamma}_{1t} \in \mathbb{R}^N$, $\boldsymbol{\mu}_{0t}, \boldsymbol{\gamma}_{0t} \in \mathbb{R}^r$, $\boldsymbol{\Sigma}_{1t} \in \mathbb{R}^{N\times N}$, $\boldsymbol{\Sigma}_{0t} \in \mathbb{R}^{r\times r}$, $\boldsymbol{\Sigma}_{10t} \in \mathbb{R}^{N\times r}$ and $\boldsymbol{\Sigma}_{01t} \in \mathbb{R}^{r\times N}$. The subscript 1 indicates observed locations and 0 is the location(s) we want to predict. Finally we use the algorithm outlined in Section 3.2 to obtain MCMC samples for the posteriors of the parameters in (50). To obtain the spatial predictive distribution, using the result in Härdle and Simar (2007), we have

$$\boldsymbol{Y}_t^* | \boldsymbol{Y}_t = \boldsymbol{y}_t \sim MGH\left(\lambda', \chi', \psi', \boldsymbol{\mu}_t', \boldsymbol{\Sigma}_t', \boldsymbol{\gamma}_t'\right), \tag{51}$$

where

$$\lambda' = \lambda - \dim\left(\boldsymbol{Y}_{1t}\right)/2 \tag{52}$$

$$\boldsymbol{\mu}_t' = \boldsymbol{\mu}_{0t} + \left(\boldsymbol{y}_{1t} - \boldsymbol{\mu}_{1t}\right)\boldsymbol{\Sigma}_{1t}^{-1}\boldsymbol{\Sigma}_{10t} \tag{53}$$

$$\boldsymbol{\Sigma}_t' = \left|\boldsymbol{\Sigma}_{1t}\right|^{1/\dim(\boldsymbol{Y}_{0t})}\left(\boldsymbol{\Sigma}_{0t} - \boldsymbol{\Sigma}_{01t}\boldsymbol{\Sigma}_{0t}^{-1}\boldsymbol{\Sigma}_{10t}\right) \tag{54}$$

$$\boldsymbol{\gamma}_t' = \left(\boldsymbol{\gamma}_{0t} - \boldsymbol{\gamma}_{1t}\boldsymbol{\Sigma}_{1t}^{-1}\right)\left(\boldsymbol{\Sigma}_{0t} - \boldsymbol{\Sigma}_{01t}\boldsymbol{\Sigma}_{1t}^{-1}\right)^{-1}\boldsymbol{\Sigma}_t' \tag{55}$$

$$\chi' = \left|\boldsymbol{\Sigma}_{1t}\right|^{1/\dim(\boldsymbol{Y}_{0t})}\left(\chi + \left(\boldsymbol{y}_{1t} - \boldsymbol{\mu}_{1t}\right)\boldsymbol{\Sigma}_{1t}^{-1}\left(\boldsymbol{y}_{1t} - \boldsymbol{\mu}_{1t}\right)'\right) \tag{56}$$

$$\psi' = \left|\boldsymbol{\Sigma}_{0t}\right|^{1/N}\left(\psi + \boldsymbol{\gamma}_{0t}'\boldsymbol{\Sigma}_{0t}^{-1}\boldsymbol{\gamma}_{0t}\right)\left|\boldsymbol{\Sigma}_{1t}\right|^{1/\dim(\boldsymbol{Y}_{0t})} - \tag{57}$$

$$\left(\left(\boldsymbol{\gamma}_{0t} - \boldsymbol{\gamma}_{1t}\boldsymbol{\Sigma}_{1t}^{-1}\right)\left(\boldsymbol{\Sigma}_{0t} - \boldsymbol{\Sigma}_{01t}\boldsymbol{\Sigma}_{1t}^{-1}\right)^{-1}\right)'\boldsymbol{\Sigma}_t'\left(\boldsymbol{\gamma}_{0t} - \boldsymbol{\gamma}_{1t}\boldsymbol{\Sigma}_{1t}^{-1}\right)\left(\boldsymbol{\Sigma}_{0t} - \boldsymbol{\Sigma}_{01t}\boldsymbol{\Sigma}_{1t}^{-1}\right)^{-1}. \tag{58}$$

We also used this result for (34).

3.4 Temporal predictive distribution

In this subsection we briefly describe one-step ahead temporal prediction. To obtain one-step ahead temporal predictions, or forecasts, we consider two particular cases; forecasting in the observed spatial location s_i, $i = 1, \ldots, N$, and in out-of-sample spatial location s_0.

Suppose we want to forecast at time $T+1$ at location s_i denoted by $D_{T+1}(s_i)$. We can easily obtain the posterior temporal predictive distribution of $D_{T+1}(s_i)$ from that of $Y_{T+1}(s_i)$, which can be calculated via

$$Y_{T+1}(s_i) = \beta_0 + \sum_{k=1}^{M} \beta_k X_{k(T+1)}(s_i) + V_{T+1}(s_i), \tag{59}$$

where

$$\begin{aligned} &V_{T+1}(s_i) = bV_T(s_i) + U_{T+1}(s_i) \quad \text{and} \\ &U_{T+1}(s_i) = \sum_{j=1}^{N} \sigma_{T+1}(s_i, s_j) Z_{T+1}(s_j) \quad \text{with} \quad \sigma^2_{T+1}(s_i, s_i) = \omega + \alpha U^2_T(s_i). \end{aligned} \tag{60}$$

We note that the correlation structure of the error term U and the distributional property of Z are independent of time.

For forecasting in a new location s_0 at time $T+1$ we define the observation $D_{T+1}(s_0)$. We already obtain the forecast values $D_{T+1}(s_i)$ and hence to obtain the forecast distribution in location s_0, we use the spatial predictive distribution described in the previous section.

4 Modelling results

The WA winter daily rainfall data described in Section 2 are analysed using our Bayesian hierarchical model. For comparision, two models are fitted: one for winter daily precipitation in years 1988-2008, and the other for winter weekly precipitation totals in year 1958-2008. Both models are fitted with 10,000 MCMC iterations, and the first 3,000 samples are discarded as burn-in. The MCMC chains converged quickly within a few hundred iterations. For brevity, these results and other MCMC diagnostics (see, Gelman and Rubin (1992)) are omitted.

4.1 Model-based analysis

Table 2 shows estimates of the parameters of the daily precipitation model. The autoregressive parameter b for the errors show a significant positive autocorrelation between precipitation on successive days. A similar positive autocorrelation is shown for the variance model. Summary statistics for the three GH parameters are also shown below. A small positive value for m_τ reflects the fact that the data is substaintially right-skewed. For simplicity, we have used the exponetional covariance function in our model so the smoothing parameter $\phi = 0.5$. The median of the spatial decay parameter θ is estimated to be 1.03, and the 95% CI varies from 0.97 to 1.09. This corresponds to an approximate 291 km spatial effective range of dependency over the study region. Summary statistics for the global effect of the four covariates are also shown in Table 2. As expected, elevation shows a negative effect on precipitation for our study region. We also find that the global effect of the covariate NINO 3.4 is positive and statistically significant, consistent with results in Crimp et al. (2014).

Parameters	Median	Std.dev	2.5%	97.5%
b (AR)	5.31×10^{-2}	2.44×10^{-2}	9.52×10^{-3}	1.05×10^{-1}
ω (ARCH)	131.81	14.56	110.88	165.80
α (ARCH)	1.44×10^{-2}	6.27×10^{-3}	3.84×10^{-3}	2.90×10^{-2}
ψ (Shape)	5.99×10^{-3}	1.15×10^{-3}	4.15×10^{-3}	8.62×10^{-3}
λ (Subclass)	6.25×10^{-2}	6.26×10^{-2}	-6.61×10^{-2}	1.76×10^{-1}
m_τ (Skewness)	3.67×10^{-4}	5.03×10^{-4}	1.49×10^{-5}	1.87×10^{-3}
θ (Spatial decay)	1.03	3.03×10^{-2}	9.73×10^{-1}	1.09
β_0 (Global intercept)	1.09	7.79×10^{-2}	9.45×10^{-1}	1.25
β_1 (NINO 3.4)	7.61×10^{-2}	5.73×10^{-2}	3.44×10^{-3}	1.86×10^{-1}
β_2 (SAMI)	-1.19×10^{-1}	3.66×10^{-2}	-1.57×10^{-1}	-7.89×10^{-2}
β_3 (IOD)	3.11×10^{-1}	5.28×10^{-1}	-2.53×10^{-1}	6.94×10^{-1}
β_4 (ELE)	-3.14×10^{-2}	4.63×10^{-3}	-1.21×10^{-1}	1.06×10^{-2}

Table 2: Posterior mean and corresponding 95% CIs of the parameters of the spatio-temporal model for the study region.

For brevity, results for the weekly precipitation model is not presented here. The effect of the four covariates on weekly precipitation totals is similar to those in Table 2. However, other parameter estimates vary considerably reflecting different skewness, temporal and spatial dependencies in the weekly data. For instance, the median of the spatial decay parameter θ is estimated to be 0.43, and the 95% CI varies from 0.28 to 0.56, corresponding to a much larger spatial effective range of dependency (697 km) in weekly precipitation totals over the study region.

4.2 Sensitivity analysis

One of the disadvantages of Bayesian analysis is due to the fact that the prior distribution for the parameters can have a significant impact on the posterior distribution, and consequently, lead to biased results. We checked the sensitivity of the model by using diffierent hyper-parameters of the Normal and truncated Normal priors. In the case where the original priors are uniform distributions, Beta distributions on the same support are used as alternatives. The results showed that our model is not very sensitive to the choice of the hyper-parameter values. For brevity these results have been omitted from the paper.

4.3 Spatial predictions of daily and weekly precipitation and its extremes

To check and validate the model's performances, we produce spatial predictions of precipitation return periods at out-of-sample locations, and compare them to empirical return preiods. As noted previously, the calculation of return period assumes that the probability of the event occurring does not vary over time and is independent of past event. Although this is only true for large precipitation, we still calculate it for all observed precipitation amounts solely for model validation purposes.

To produce out-of-sample spatial predictions for return periods, we hold back one weather station at a time and use the previously outlined algorithm and (51) to obtain samples from the spatial predictive distribution for that location, then calculate return periods using the method described in Section 1.1. Finally the 95% prediction credible interval (CI) for the return period curve is produced based on 1,000 simulations. The left panel of Figure 2 shows the spatial prediction of daily precipitation return periods at Site 9619. The red dots are empirical return periods, whereas the dashed lines represent the 95% CI. We can also compare samples from the spatial predictive distribution directly to observed daily precipitation through a Q-Q plot as shown in the right panel of Figure 2. Observed quantiles of winter daily precipitation observed in year 1988-2008 are on the x-axis, and mean simulated quantiles are on the y-axis. The same two

graphs for winter weekly precipitation totals in years 1958-2008 are also produced at Site 9619. The model appears to produce excellent spatial predictions, including for extreme precipitation events, and accommodate different shapes and skewness. It is important to recognise that the CIs for return periods constructed under the proposed model are much narrower than those constructed under extreme value theory, e.g. Li et al. (2005). Out-of-sample predictions of return periods and the Q-Q plots for both daily and weekly precipitation data at all other weather stations are presented in Figure 7, 8, 9 and 10 in the Appendix. Similar conclusions can be drawn from those plots. Given the demanding criteria of accurate out-of-sample prediction of extreme precipitation, these results are quite impressive. However, in Figure 2, the out-of-sample spatial prediction for winter daily precipitation return periods shows a noticible misalignment at 0 mm and for very small precipitation amounts. This might be due to model error or measurement errors of small precipitation amounts.

Figure 3 shows the empirical and modelled spatio-temporal variograms for both daily and weekly

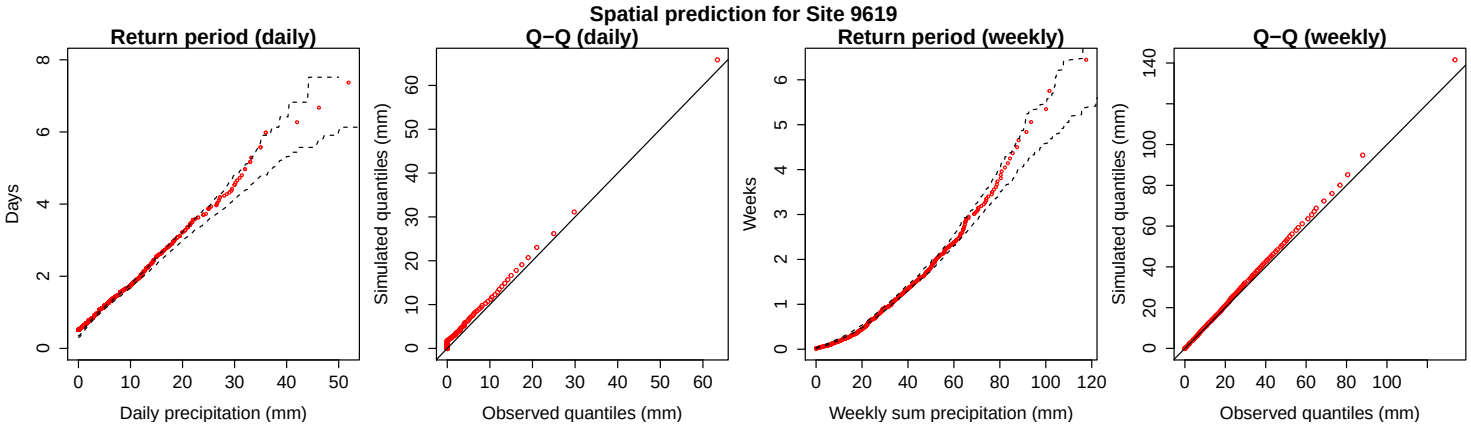

Figure 2: Out-of-sample spatial prediction of daily and weekly sum precipitation for Site 9619, based on 1,000 simulations.

precipitation data. For both daily and weekly data, the modelled variogram closely matches the empirical variogram. Also it is worth noting that total weekly precipitation shows much less temporal dependency than daily precipitation. Furthermore, there is an increase in the variogram for lag 1 and greater for distance close to zero, which is expected for rainfall data.

Another frequently used uncertainty measure for precipitation and its extremes is precipitation duration. Figure 4 shows out-of-sample spatial predictions of mean duration (days) over/below thresholds for winter daily precipitation. The box-plots are based on 1,000 simulations, and the red dots represent observed mean durations (days) over/below thresholds for 1988-2008. It shows that the model matches observations at most out-of-sample locations, especially for duration over large thresholds, indicating that the model efficiently represents spatial and short-term temporal dependencies. Similarly, out-of-sample spatial predictions of numbers of weeks over/below thresholds are presented in Figure 5, with the top panel showing spatial predictions of the number of weeks with no rainfall at each out-of-sample weather station. The red dots represent observed numbers of winter weeks over/below thresholds for 1958-2008. Again the out-of-sample predictions are in close agreement with the observed data.

Finally, we produce predictive maps for numbers of weeks over/below thresholds. For prediction purposes, we define 150 grid points over our study region and obtained interpolated predicted values of the winter weekly precipitation totals using data from the 21 weather stations for the time period 1958-2008. Figure 6 shows the predicted mean spatial pattern of the number of weeks with aggregated precipitation greater than 90 mm and exactly 0 mm, respectively. The values that appear on the map represent the observed values. The plots are accompanied by their standard deviation plots. The model shows close

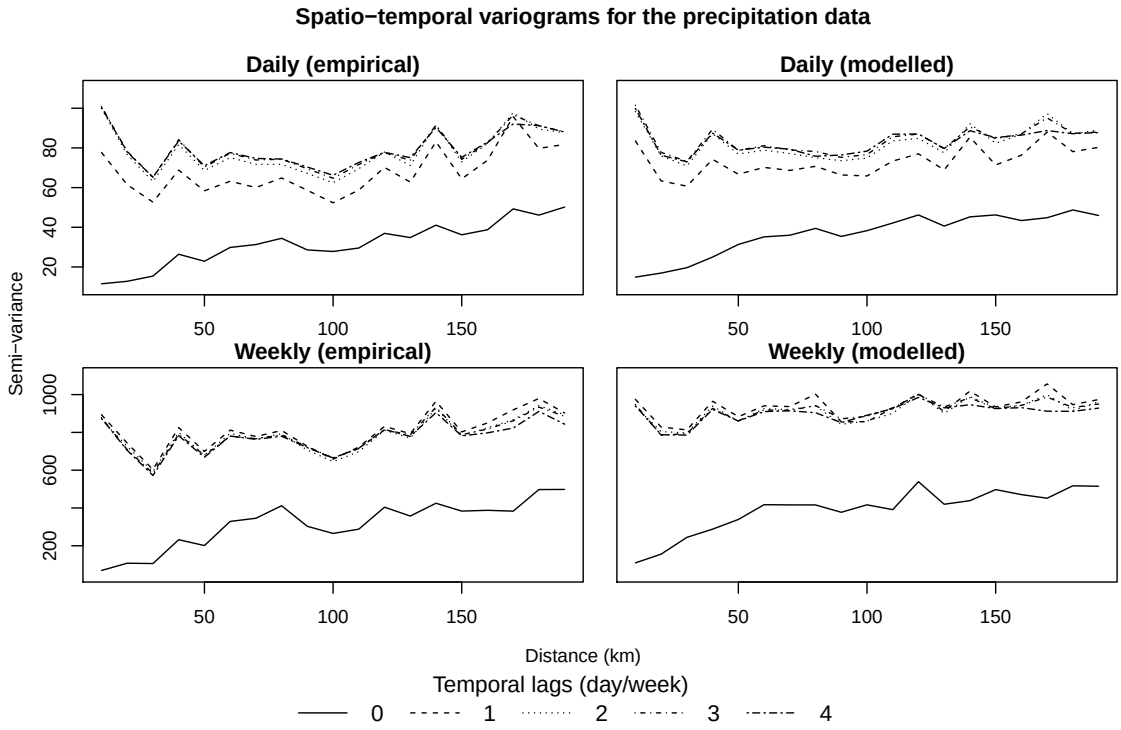

Figure 3: Empirical and modelled spatio-temporal variograms for both daily and weekly precipitation data. Modelled variograms are based on 1,000 simulations.

agreement with observations across most of the study region.

5 Discussion

In this paper, we developed a Bayesian heirarchical model that utilises the generalised hyperbolic process for producing spatial predictions and measures of uncertainty for spatio-temporal data that is heavy-tailed and subject to substantial and varying skewness. Unlike models based on extreme value theory, which only model maxima of finite-sized blocks or exceedances above a large threshold, the proposed model uses all the data available efficiently, and hence not only fits the extremes but also models the entire rainfall distribution. We applied the method to both winter daily precipitation and weekly precipitation totals across a study region in south-west Western Australia. Our example shows that the proposed model can accommodate spatially and temporally varying volatility and skewness, and efficiently represents spatial and short-term temporal dependencies.

In future research we plan to extend the method to more general treatments of similar spatio-temporal processes in different ways. Firstly the coefficients β_k can be made spatially varying and temporally dynamic as in Bakar et al. (2014). Spatially varying coefficient models are often used to address the point-to-area problem of covariates which only vary over time, but whose impact may vary across space (Gelfand et al., 2003), whereas dynamic linear models (West et al., 1985; Stroud et al., 2001) are popular for modelling data with seasonal variations. Recent applications include Dou et al. (2010); Ghosh et al. (2010); Mahmoudian and Mohammadzadeh (2014) and Sahu and Bakar (2012). Secondly the autoregressive parameter b can be generalised in a similar way (Bakar et al., 2014; Cressie and Wikle, 2011).

In this research, we adopted a censoring approach for modelling zero precipitation. Similiar approaches were also used in Stein (1992), Glasbey and Nevison (1997) and Yuen and Guttorp (2014). An alternative

approach for modelling zero precipitation considers a mixed probability distribution composed of a discrete at zero and a continuous distribution for non-zero precipitation; see e.g. Srikanthan and McMahon (1999) and Woolhiser (1992). Another approach involves using the empirical distribution below a small threshold; see Dupuis (2012). It is also important to recognise that when analysing weekly precipitation totals, the data became less zero-inflated, and the censoring approach coupled with the flexibility of the generalised hyperbolic structure was able to model zero precipitation very well. Nevertheless, the proposed model could be adapted for the above alternative methods for modelling zero precipitation.

There are some limitations of the proposed method. First of all, the model will fail when applying to massive spatio-temporal dataset due to the big-n problem (Cressie and Wikle, 2011). Second, in our research the temporal non-stationarity is modelled through heteroscedasticity but we assumed the spatial process is stationary, which might not be true for large study regions. However, complicated spatial and temporal interactions in the model make it difficult to generalise it to spatially non-stationary processes.

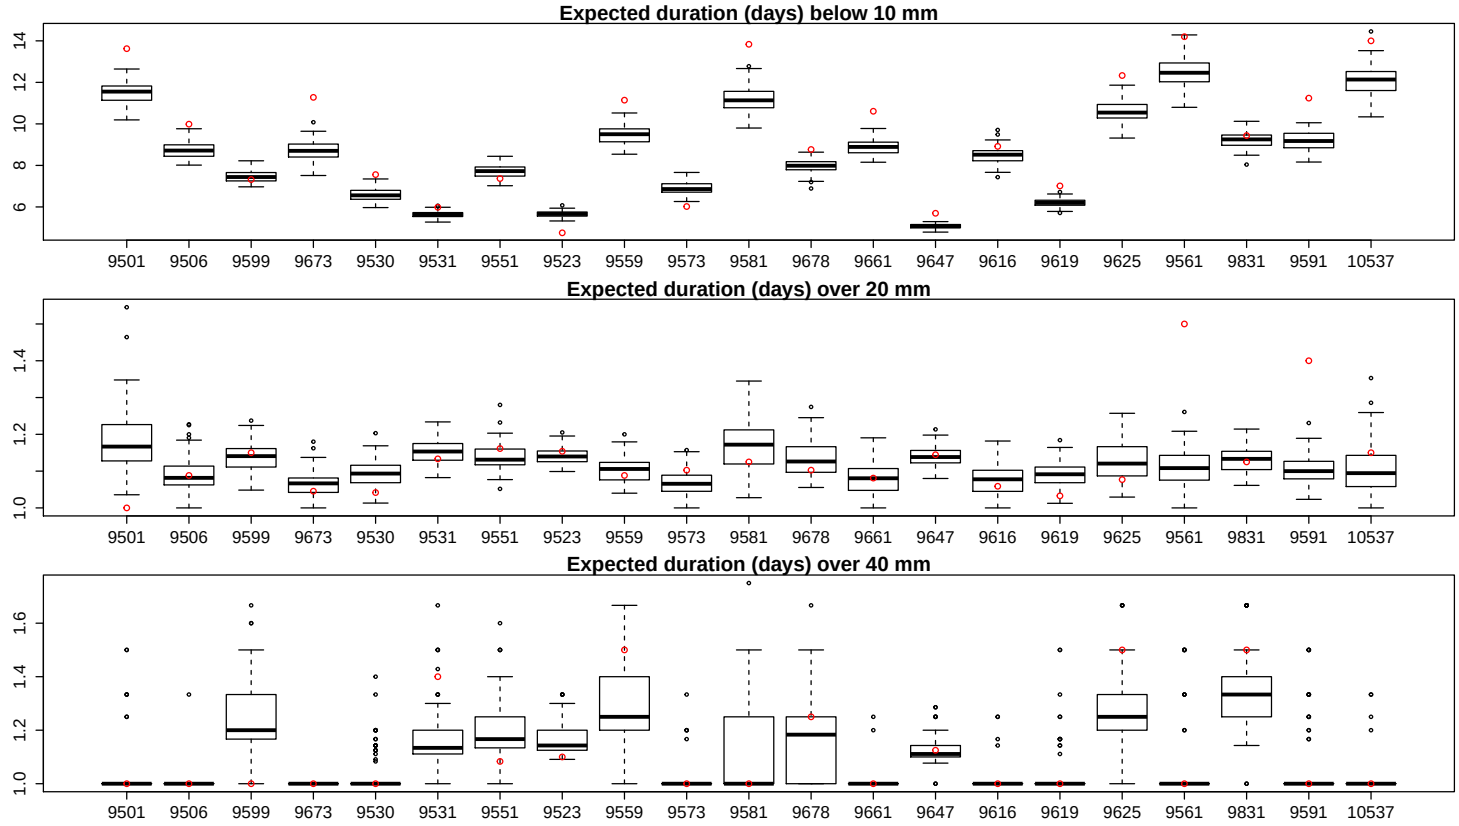

Figure 4: Out-of-sample spatial predictions of expected durations (days) over/below thresholds based on the proposed hierarchical model and 1,000 simulations. The red dots are the mean duration over/below thresholds observed in year 1988-2008.

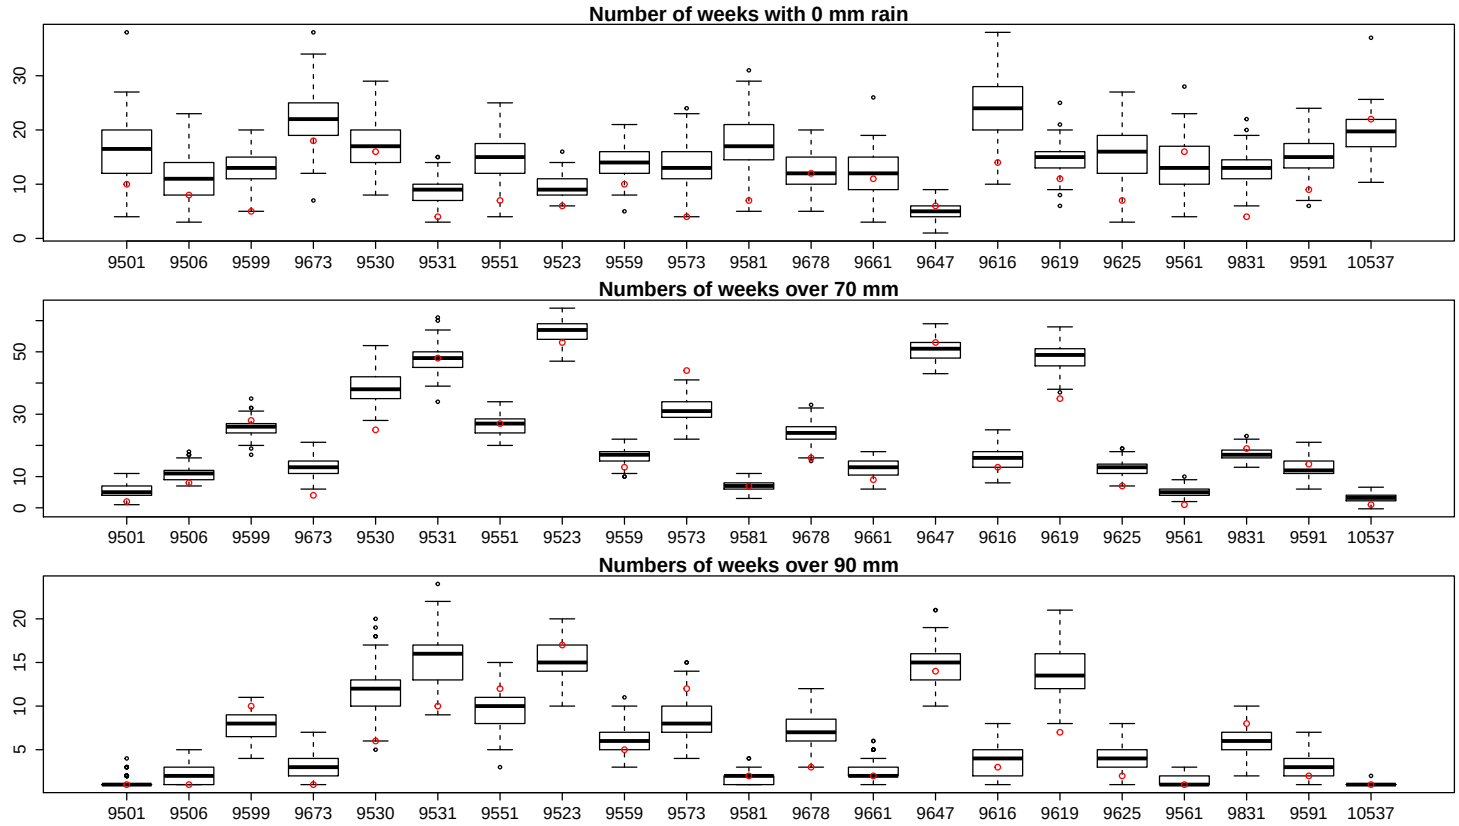

Figure 5: Out-of-sample spatial predictions of numbers of weeks over/below thresholds based on the proposed hierarchical model and 1,000 simulations. The red dots are the numbers of weeks over/below thresholds observed in year 1958-2008.

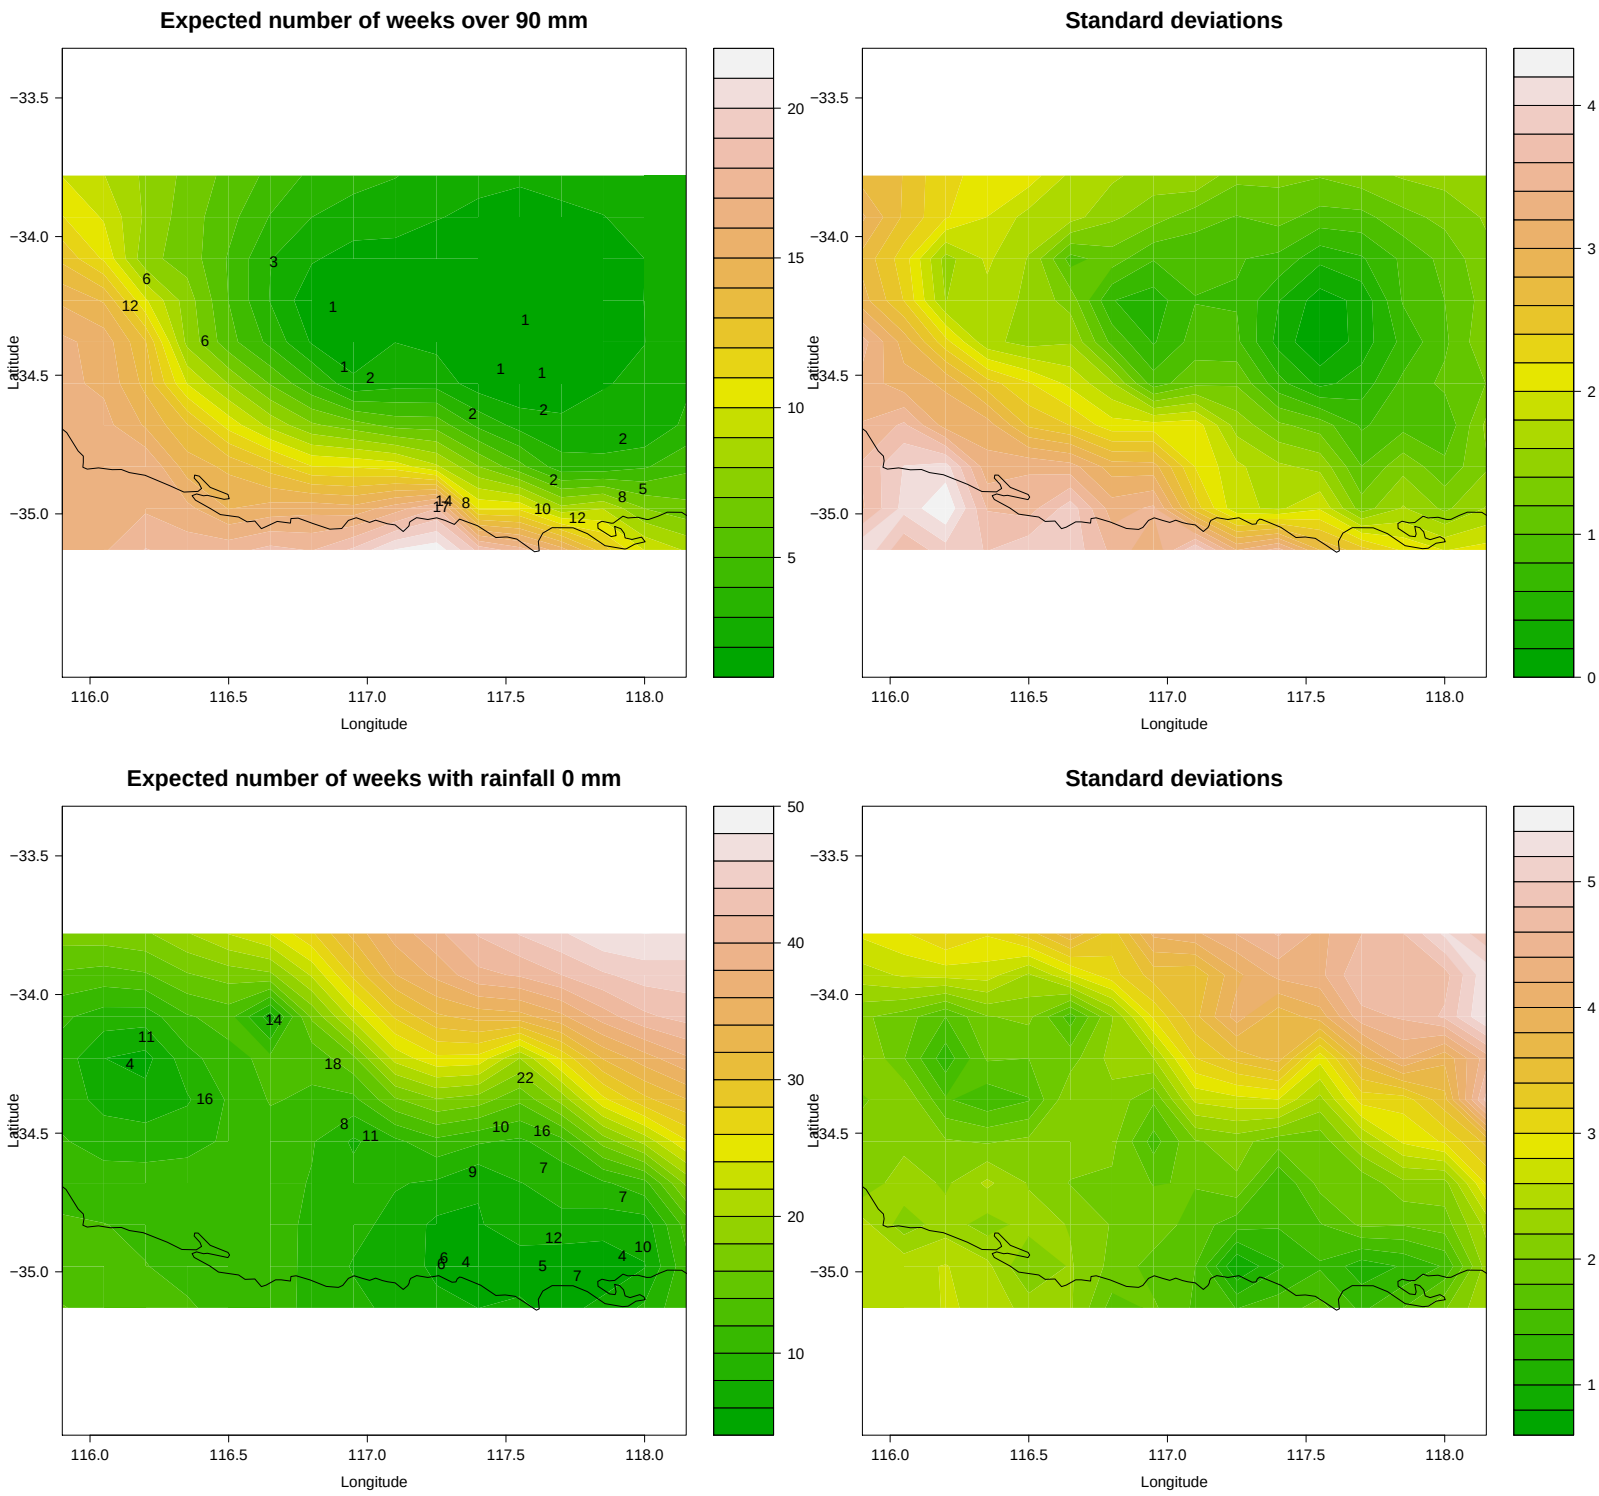

Figure 6: Predicted expected numbers of weeks with >90 mm and 0 mm rain and the corresponding standard deviations for the study region, based on the proposed hierarchical model and 1,000 simulations. The numbers represent the numbers of weeks with >90 mm and 0 mm rain observed in winter weekly precipitation sum in year 1958-2008.

Appendix

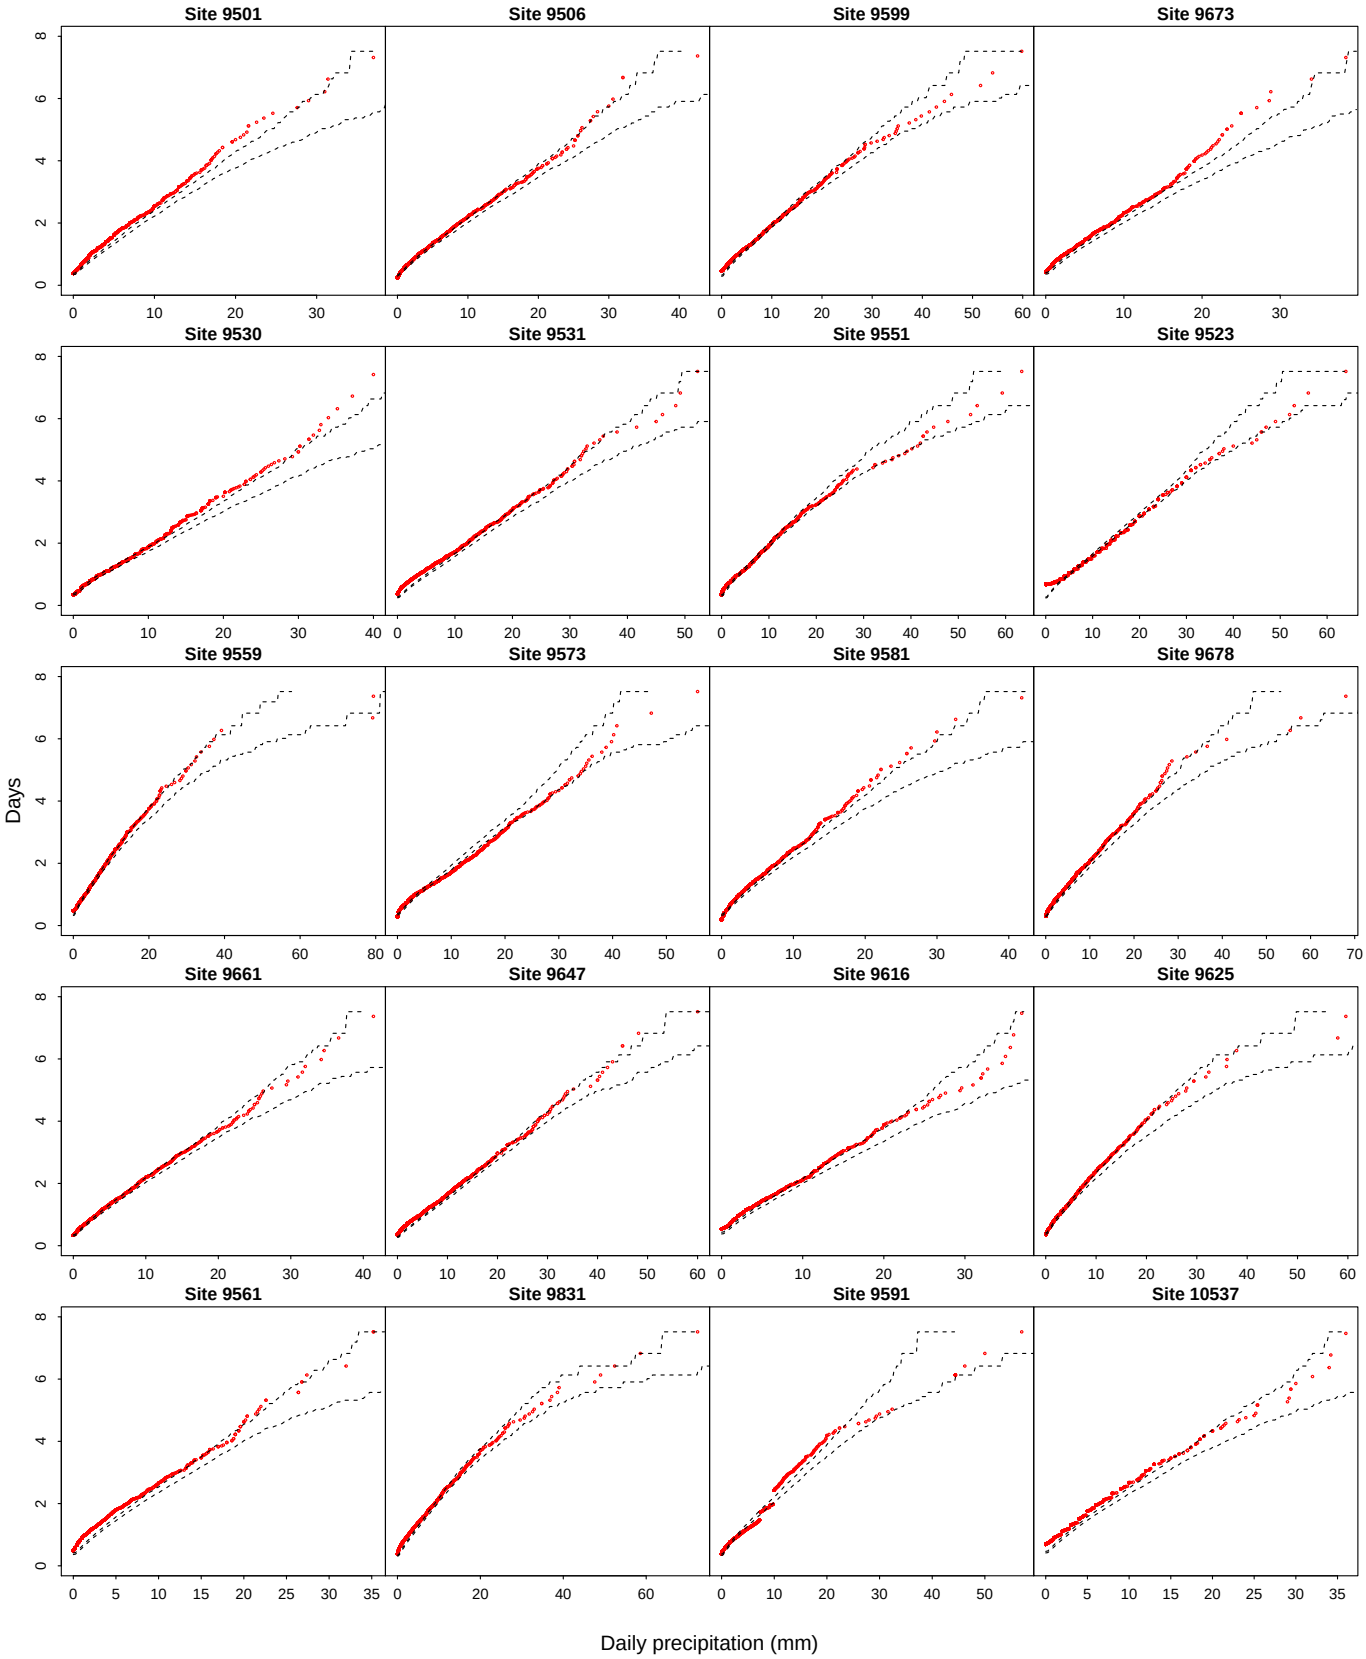

Figure 7: Out-of-sample spatial predictions of precipitation return periods (days). The dashed lines are the 95% CIs, which are based on the proposed hierarchical model and 1,000 simulations; the red dots are the observed precipitation return periods (days) of winter daily precipitation observed in year 1988-2008.

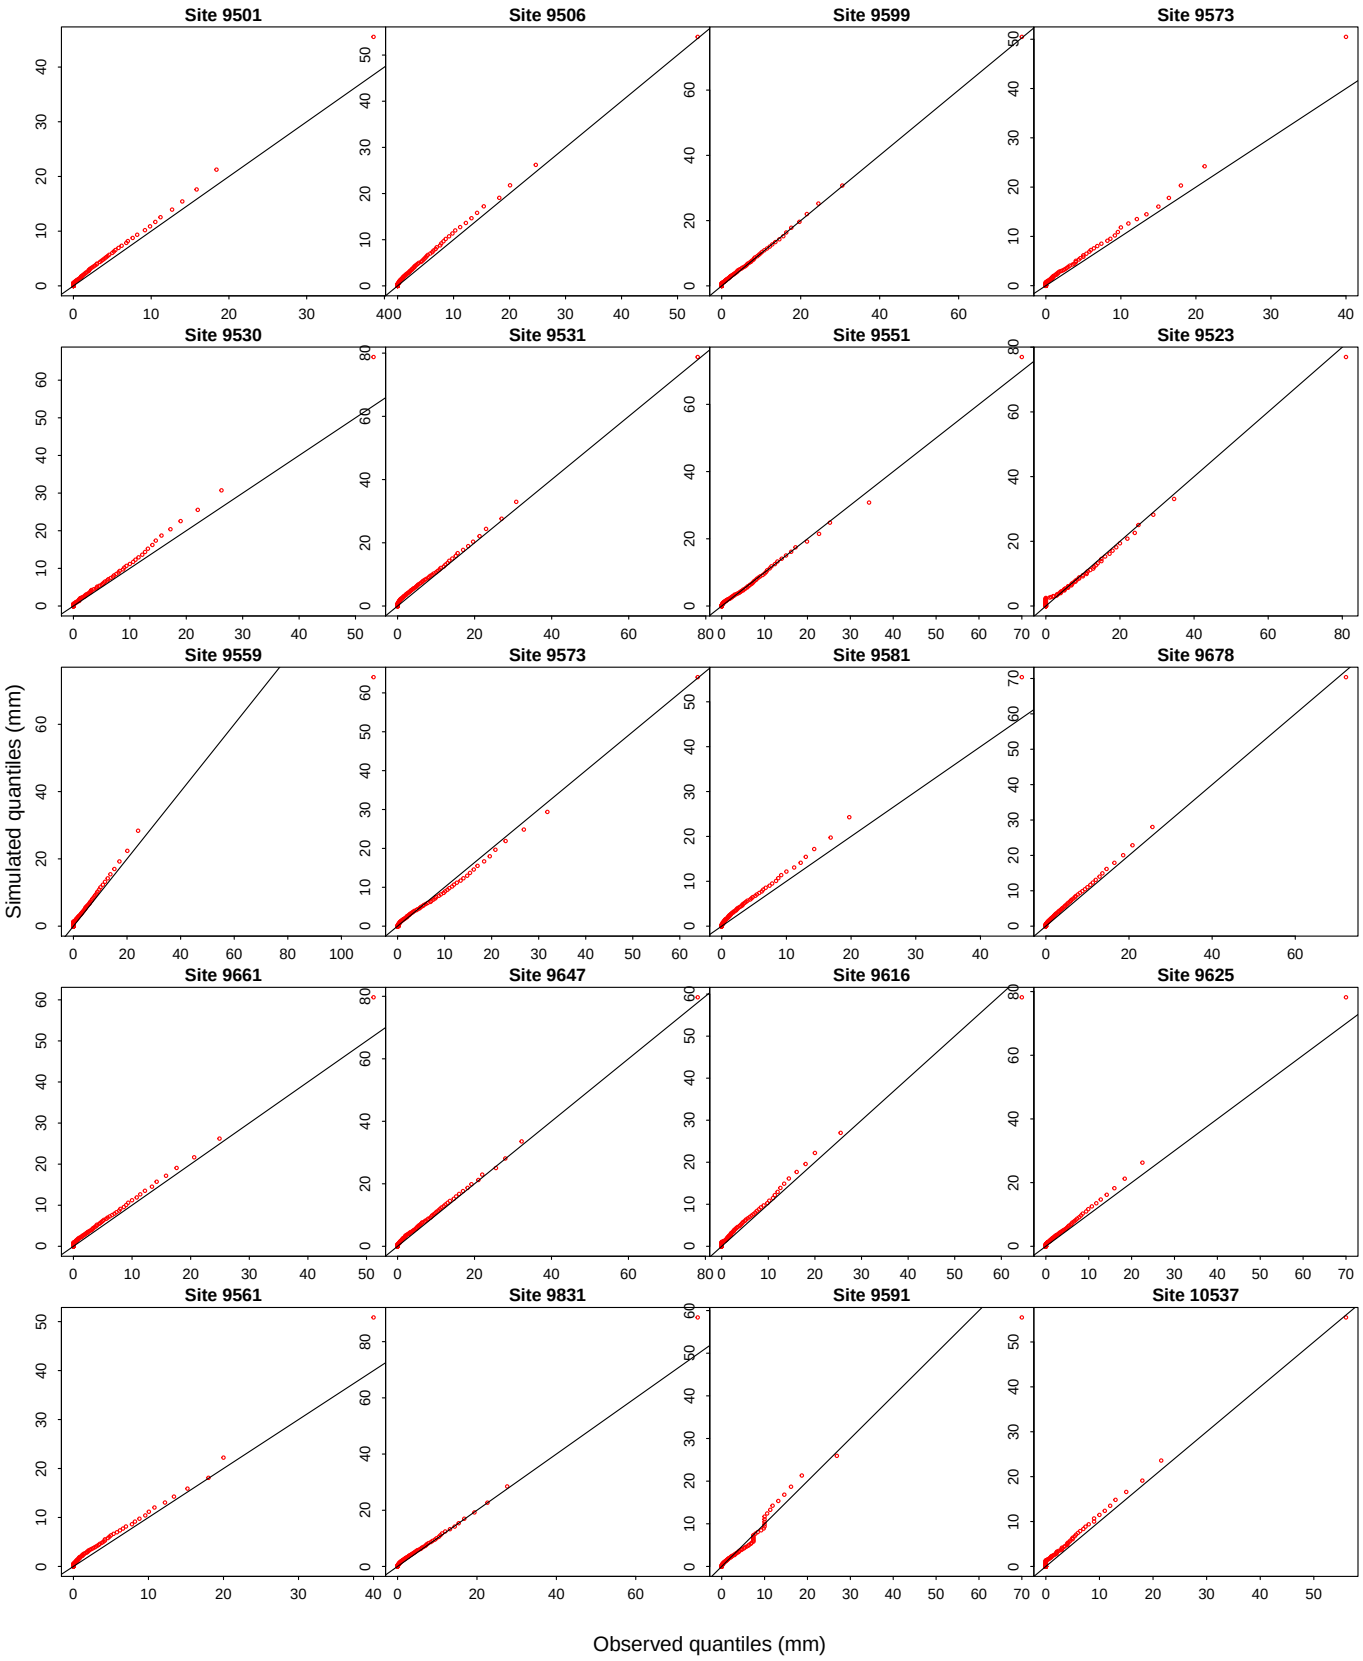

Figure 8: Out-of-sample spatial predictions of distributions of daily precipitation based on the proposed hierarchical model and 1,000 simulations. Observed quantiles of winter daily precipitation observed in year 1988-2008 are on the x-axis, and mean simulated quantiles are on the y-axis.

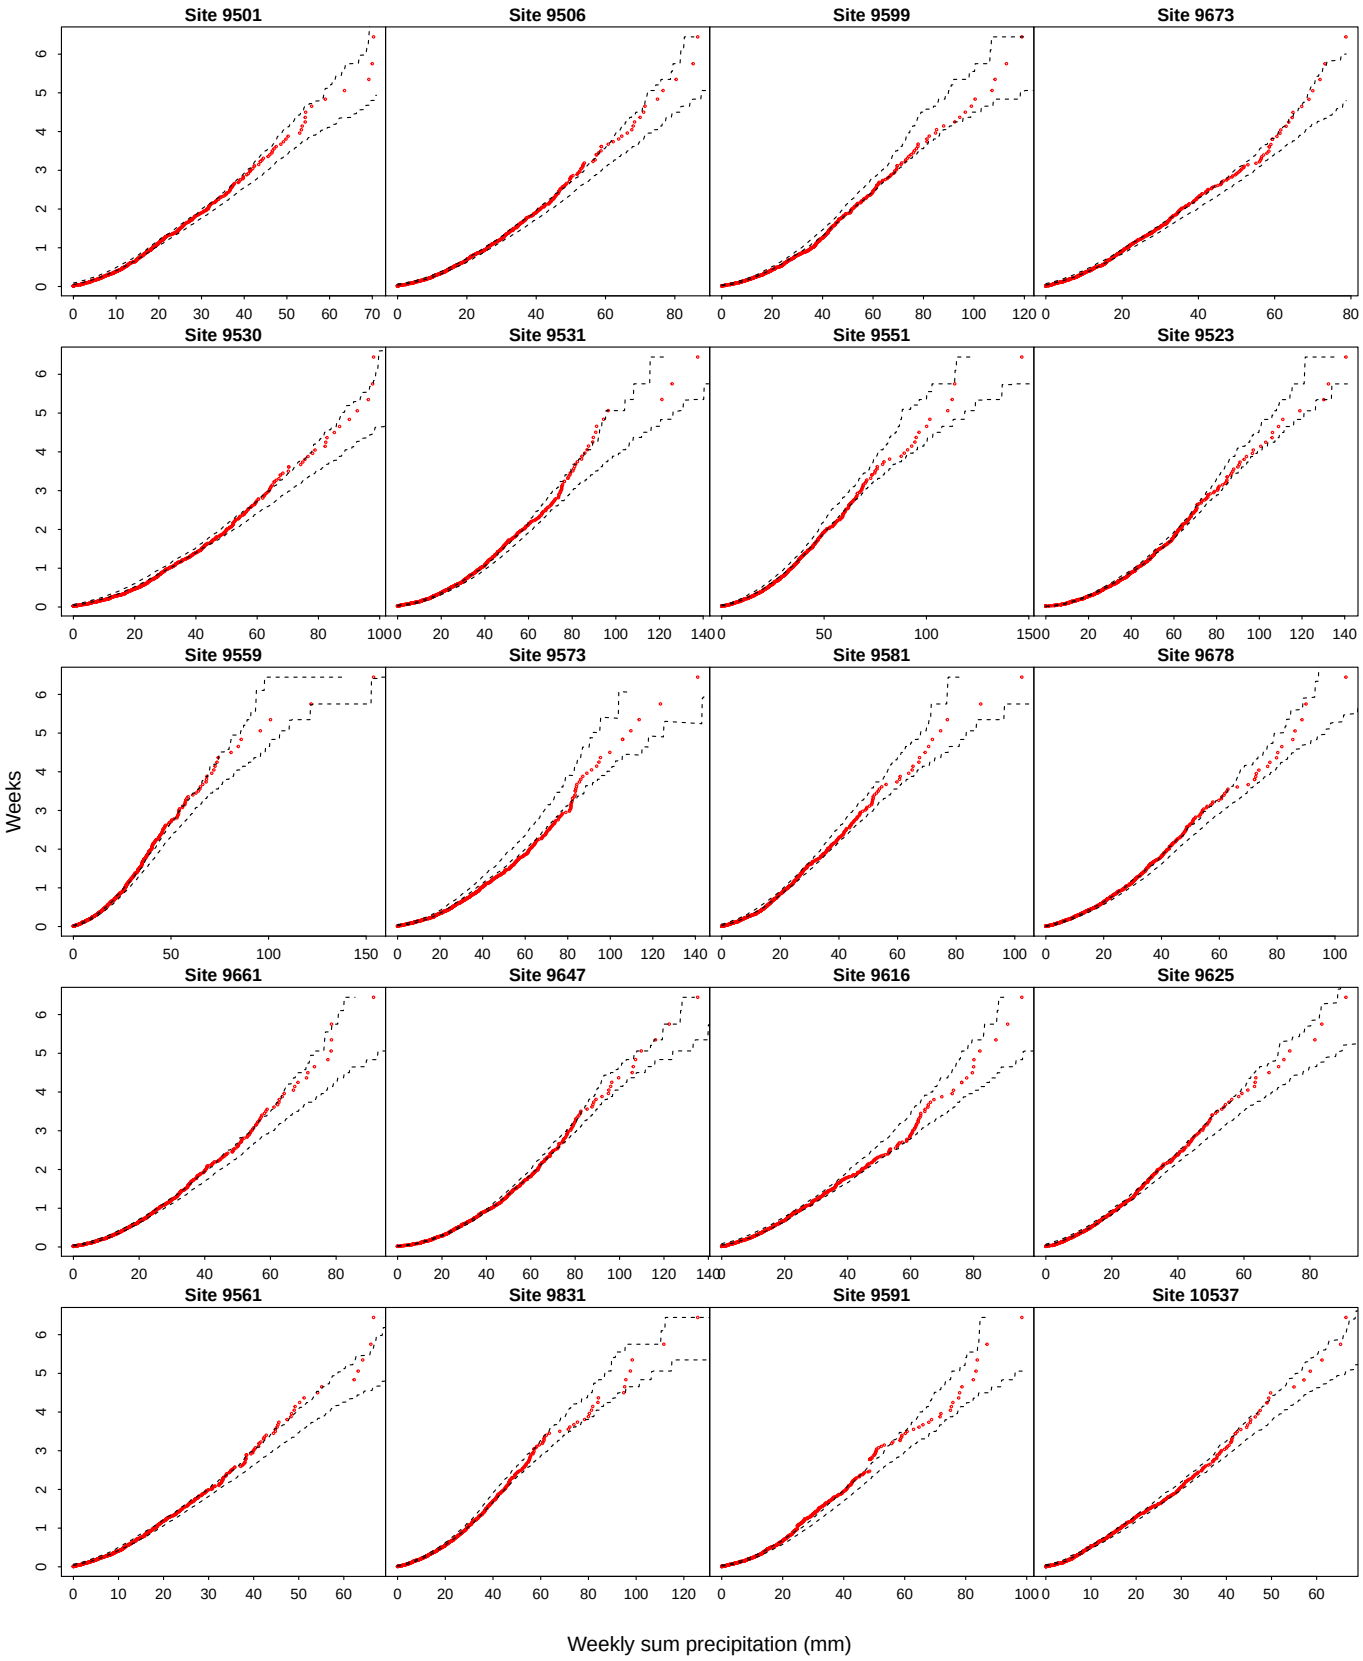

Figure 9: Out-of-sample spatial predictions of precipitation return periods (weeks). The dashed lines are the 95% CIs, which are based on the proposed hierarchical model and 1,000 simulations; the red dots are the observed precipitation return periods (days) of winter weekly precipitation sum observed in year 1958-2008.

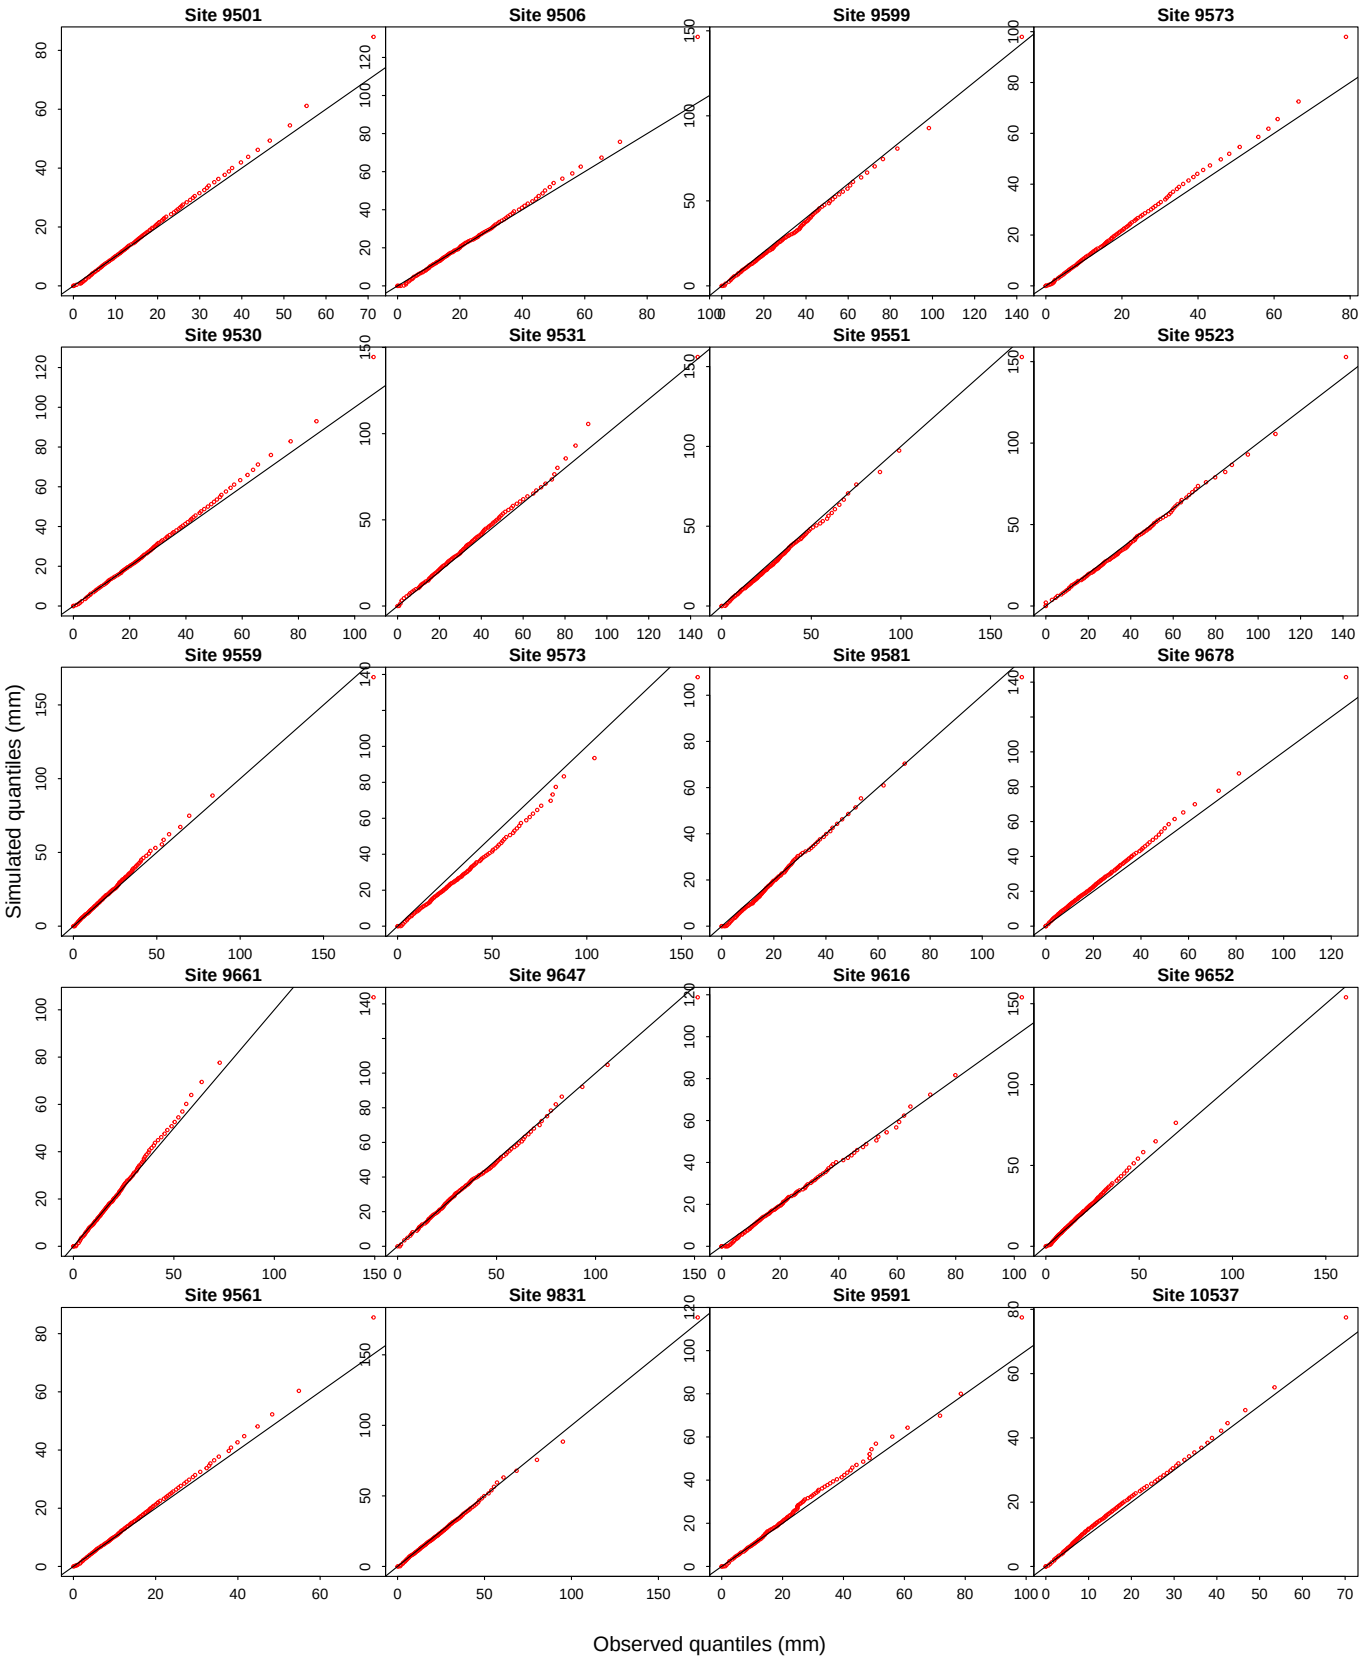

Figure 10: Out-of-sample spatial predictions of distributions of weekly precipitation sum based on the proposed hierarchical model and 1,000 simulations. Observed quantiles of winter weekly precipitation sum observed in year 1958-2008 are on the x-axis, and mean simulated quantiles are on the y-axis.

References

Ashcroft, M., K. French, and L. Chisholm (2011). An evaluation of environmental factors affecting species distributions. *Ecological Modelling 222*, 524–531.

Ashcroft, M. and J. Gollan (2012). Fine-resolution (25 m) topoclimatic grids of near-surface (5 cm) extreme temperatures and humidities across various habitats in a large (200 - 300 km) and diverse region. *International Journal of Climatology 32*, 2134–2148.

Bakar, K., P. Kokic, and H. Jin (2014). A spatio-dynamic model for assessing frost risk in south-east Australia (Under review).

Banerjee, S., A. E. Gelfand, and B. P. Carlin (2004). *Hierarchical modeling and analysis for spatial data*. Crc Press.

Barndorff-Nielsen, O. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences 353*(1674), 401–419.

Barndorff-Nielsen, O. E. and N. Shephard (2001). *Normal modified stable processes*. MaPhySto, Department of Mathematical Sciences, University of Aarhus.

Coles, S. G. and J. A. Tawn (1996). A bayesian analysis of extreme rainfall data. *Applied Statistics*, 463–478.

Cooley, D., D. Nychka, and P. Naveau (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association 102*(479), 824–840.

Cressie, N. and C. K. Wikle (2011). *Statistics for spatio-temporal data*. John Wiley & Sons.

Crimp, S., K. S. Bakar, P. Kokic, H. Jin, N. Nicholls, and M. Howden (2014). Bayesian space–time model to analyse frost risk for agriculture in southeast Australia. *International Journal of Climatology*.

De Oliveira, V. (2005). Bayesian inference and prediction of gaussian random fields based on censored data. *Journal of Computational and Graphical Statistics 14*(1).

Dou, Y., N. D. Le, J. V. Zidek, et al. (2010). Modeling hourly ozone concentration fields. *The Annals of Applied Statistics 4*(3), 1183–1213.

Dupuis, D. J. (2012). Modeling waves of extreme temperature: the changing tails of four cities. *Journal of the American Statistical Association 107*(497), 24–39.

Eberlein, E. and E. A. v. Hammerstein (2004). Generalized hyperbolic and inverse gaussian distributions: limiting cases and approximation of processes. In *Seminar on Stochastic Analysis, Random Fields and Applications IV*, pp. 221–264. Springer.

Feng, S., S. Nadarajah, and Q. Hu (2007). Modeling annual extreme precipitation in China using the generalized extreme value distribution. *85*(5), 599–613.

Fisher, R. A. and L. H. C. Tippett (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Volume 24, pp. 180–190. Cambridge Univ Press.

Fowler, H. J., S. Blenkinsop, and C. Tebaldi (2007). Linking climate change modelling to impacts studies: recent advances in downscaling techniques for hydrological modelling. *International Journal of Climatology 27*, 1547–1578.

Gaetan, C. and M. Grigoletto (2007). A hierarchical model for the analysis of spatial rainfall extremes. *Journal of agricultural, biological, and environmental statistics 12*(4), 434–449.

Gelfand, A. E., H.-J. Kim, C. Sirmans, and S. Banerjee (2003). Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association 98*(462), 387–396.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013). *Bayesian data analysis*. CRC press.

Gelman, A. and D. B. Rubin (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, 457–472.

Ghosh, S. K., P. V. Bhave, J. M. Davis, and H. Lee (2010). Spatio-temporal analysis of total nitrate concentrations using dynamic statistical models. *Journal of the American Statistical Association 105*(490), 538–551.

Glasbey, C. and I. Nevison (1997). Rainfall modelling using a latent Gaussian variable. In *Modelling Longitudinal and Spatially Correlated Data*, pp. 233–242. Springer.

Gnedenko, B. (1943). Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of mathematics*, 423–453.

Härdle, W. and L. Simar (2007). *Applied multivariate statistical analysis*, Volume 2/2007. Springer.

Keating, B. A., P. S. Carberry, G. L. Hammer, M. E. Probert, M. J. Robertson, D. Holzworth, N. I. Huth, J. N. Hargreaves, H. Meinke, Z. Hochman, and et al. (2003). An overview of APSIM, a model designed for farming systems simulation. *European Journal of Agronomy 18*, 267288.

Kokic, P., H. Jin, and S. Crimp (2013). Improved point scale climate projections using a block bootstrap simulation and quantile matching method. *Climate Dynamics 41*, 853–866.

Li, Y., W. Cai, and E. Campbell (2005). Statistical modeling of extreme rainfall in southwest Western Australia. *Journal of Climate 18*(6).

Liu, Y., P. Kokic, and K. S. Bakar (2014). Censored bayesian hierarchical modelling of extreme precipitation. *Environmetrics*.

Love, S. (2012). *Spatial Modelling of Extreme Rainfall*. Ph. D. thesis, University of New South Wales.

Mahmoudian, B. and M. Mohammadzadeh (2014). A spatio-temporal dynamic regression model for extreme wind speeds. *Extremes 17*(2), 221–245.

Matérn, B. (1986). *Spatial Variation (Second ed.)*. Springer-Verlag.

Mencia, J. and E. Sentana (2004). Estimation and testing of dynamic models with generalised hyperbolic innovations.

Nadarajah, S. and D. Choi (2007). Maximum daily rainfall in South Korea. *Journal of Earth System Science 116*(4), 311–320.

Nan, S. and J. Li (2003). The relationship between the summer precipitation in the Yangtze river valley and the boreal spring southern hemisphere annular mode. *Geophysical Research Letters 30*(24).

Pickands III, J. (1975). Statistical inference using extreme order statistics. *the Annals of Statistics*, 119–131.

Ruprecht, J., Y. Li, E. Campbell, and P. Hope (2005). How extreme south-west rainfalls have changed.

Sahu, S. K. and K. Bakar (2012). A comparison of Bayesian models for daily ozone concentration levels. *Statistical Methodology 9*(1), 144–157.

Saji, N., B. N. Goswami, P. Vinayachandran, and T. Yamagata (1999). A dipole mode in the tropical Indian Ocean. *Nature 401*(6751), 360–363.

Schepen, A., Q. Wang, and D. Robertson (2012). Evidence for using lagged climate indices to forecast Australian seasonal rainfall. *Journal of Climate 25*(4), 1230–1246.

Shao, Q., Q. Wang, and L. Zhang (2013). A stochastic weather generation method for temporal precipitation simulation. In *20th International Congress on Modelling and Simulation, Adelaide, Australia, 16 December, 2013*, pp. 2681–2687.

Srikanthan, R. and T. McMahon (1999). Stochastic generation of annual, monthly and daily climate data: A review. *Hydrology and Earth System Sciences 5*(4), 653–670.

Stein, M. L. (1992). Prediction and inference for truncated spatial data. *Journal of Computational and Graphical Statistics 1*(1), 91–110.

Stroud, J. R., P. Müller, and B. Sansó (2001). Dynamic models for spatiotemporal data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*(4), 673–689.

Tetzlaff, D., S. Uhlenbrook, and P. Molnar (2005). Significance of spatial variability in precipitation for process-oriented modelling: results from two nested catchments using radar and ground station data. *Hydrology & Earth System Sciences 9*.

Tierney, L. (1994). Markov chains for exploring posterior distributions. *The Annals of Statistics*, 1701–1728.

Turkman, K. F., M. A. Turkman, and J. Pereira (2010). Asymptotic models and inference for extremes of spatio-temporal data. *Extremes 13*(4), 375–397.

Von Mises, R. (1936). La distribution de la plus grande de n valeurs. *Rev. math. Union interbalcanique 1*(1).

West, M., P. J. Harrison, and H. S. Migon (1985). Dynamic generalized linear models and Bayesian forecasting. *Journal of the American Statistical Association 80*(389), 73–83.

Woolhiser, D. (1992). Modeling daily precipitation: progress and problems. *Statistics in the Environmental and Earth Sciences 5*, 71–89.

Ye, W. and Y. Li (2011). A method of applying daily GCM outputs in assessing climate change impact on multiple day extreme precipitation for Brisbane River Catchment. In *19th international Congress on modelling and simulation*, pp. 12–16.

Yuen, R. and P. Guttorp (2014). A hierarchical Gauss-Pareto model for spatial prediction of extreme precipitation. Technical report, University of Michigan.